symmetry

MDPI

Article

From Commutation to Modal Symmetry: Half-Cycle Symmetry, Detuning Duality, and Symmetry Breaking in LLC Resonant DC–DC Converters

Nikolay Hinov [1,2]

[1] Centre of Excellence "National Center of Mechatronics and Clean Technologies", 1000 Sofia, Bulgaria; hinov@tu-sofia.bg
[2] Department of Computer Systems, Faculty of Computer Systems and Technologies, Technical University of Sofia, 1000 Sofia, Bulgaria

Abstract

Inductor–inductor–capacitor (LLC) resonant converters can produce nearly symmetric waveforms without achieving low-loss switching. This paper separates these properties, using a common framework for a resonant circuit containing two inductive elements and one capacitor. An analytical gain model is combined with a time-domain model that distinguishes positive secondary conduction, an open secondary circuit, and negative secondary conduction. The framework compares the two switching half-cycles, their conduction sequences, the gain at reciprocal normalized frequencies, and the current available for zero-voltage switching (ZVS). Finite magnetizing inductance breaks reciprocal gain symmetry, while a balanced converter can retain close half-cycle correspondence. Controlled bridge, timing, and rectifier asymmetries increase the waveform mismatch. Independent LTspice checks reproduce three representative mode families and agree with the low-order gain within 0.45–3.55%. They also confirm that a high-gain operating point can fail ZVS despite good waveform symmetry. Highlighted current paths, a datasheet-informed semiconductor-loss budget, and input/output ripple spectra connect the descriptors to practical design questions. The loss estimates are conditional analytical scenarios, and the filtering study uses a separate coupled reduced-order model. The framework supports reproducible design screening; it does not replace detailed device simulation or hardware validation.

Keywords: LLC resonant converter; half-cycle symmetry; modal symmetry; P/O/N modes; reciprocal detuning; zero-voltage switching; circulating energy; LTspice validation; reproducibility; tolerance analysis

Academic Editors: Michel Planat and Stefano Profumo

Received: 11 August 2026
Revised: 13 September 2026
Accepted: 16 September 2026
Published: 18 September 2026

1. Introduction

LLC (inductor–inductor–capacitor) resonant DC–DC converters are widely used where high efficiency, galvanic isolation, and high power density are required [1–4]. Their behavior results from the interaction of the series resonant branch L_r–C_r, transformer magnetizing inductance L_m, rectifier, output filter, and switching-frequency control [5–8]. This interaction enables primary-side zero-voltage switching (ZVS), secondary-side current relief, and voltage regulation over a finite input–load corridor, but it also creates several load-dependent conduction sequences [9–12].

Classical LLC analysis follows two complementary directions. Fundamental harmonic approximation (FHA) supplies compact gain and impedance relations for initial design, but it suppresses exact rectifier intervals and switching transitions [4,6,13–16]. Time-domain, state-plane, and operation-mode methods resolve positive-secondary-conduction (P), open-secondary (O), and negative-secondary-conduction (N) intervals more accurately [5,8–12], while design-oriented studies optimize gain range, resonant-tank current, and efficiency. These approaches are indispensable, but they normally treat gain, modal sequence, ZVS feasibility, and reactive stress as separate constraints rather than as explicitly distinguishable symmetry domains.

A separate research line has used symmetry and duality to organize resonant-converter behavior [17–19]. Those studies provide family-level or regime-level language. Related LLC trajectory and commutation analyses [20–22] and design and efficiency methods [23–27] supply topology-specific operating and optimization results. The quasi-boundary method [28] and quasi-resonant ZVS/ZCS symmetry map [29] further motivate the present descriptor framework. General resonant-converter analyses [30–33] provide the broader family context. These prior contributions do not replace the event-aligned residual, reciprocal-detuning decomposition, and combined modal/ZVS/circulation screen developed here.

The unresolved engineering gap is therefore not another exact LLC model. It is a coordinated descriptor layer that answers four different questions without conflating them: whether reciprocal detuning preserves gain, whether the positive and negative half-cycles satisfy the expected state transformation, whether their rectifier-stage words correspond, and whether the transition current has the correct magnitude and polarity for ZVS. The practical value is to rank operating points, expose why an apparently symmetric waveform may still fail commutation, quantify the reactive-energy cost of a desired mode, and select a smaller set of points for parasitic-aware simulation and hardware testing.

The central premise is that LLC operation contains several distinct, only partially coincident symmetries:

1. Half-cycle symmetry, generated by the bridge polarity reversal after $T_s/2$;
2. Frequency-domain reflection, represented by $f_n \mapsto 1/f_n$ around the series-resonance center in logarithmic coordinates;
3. Commutation symmetry, expressed through whether the resonant current supplies the charge needed to commutate the bridge-node capacitances during dead time;
4. Path symmetry, determined by the availability of primary, transformer, and rectifier current paths;
5. Modal symmetry, determined by the ordered P/O/N conduction word within each half-cycle.

These symmetries should not be conflated. An operating point may show almost exact half-cycle and modal correspondence while failing the ZVS charge condition; reciprocal-frequency partners may have different gains because finite L_m breaks that mapping; and a ZVS-capable point may incur excessive circulating energy. The useful design question is therefore not simply whether the topology is symmetric, but which symmetry is active, which mechanism breaks it, and what electrical cost accompanies the desired behavior.

The contributions of this work are as follows:

1. A topology-specific five-coordinate LLC signature covering half-cycle parity, detuning side, commutation tendency, path feasibility, and canonical P/O/N modal structure;
2. An analytical decomposition proving that the reactive FHA term is reciprocal-frequency invariant while finite L_m introduces an explicit symmetry-breaking term;

3. A Z_2 half-cycle state transformation and a canonical P/O/N modal-word comparison;
4. An event-aligned, RMS-normalized half-cycle residual, a canonical modal-word score, an oriented ZVS charge margin, and a circulating-energy cost;
5. A frozen and machine-checked numerical workflow with separate settling, time-step, event-location, threshold-sensitivity, and physical-parameter tolerance studies;
6. Independent parasitic-aware LTspice spot checks at $f_n = 0.45$, 0.80, and 1.25, including modal-family, gain, current-stress, loss, and turn-on-voltage comparisons.

The LLC converter is used here as a physical reference system for investigating symmetry and reciprocal-detuning relations, as well as their limitations. The circuit equations and conduction-state descriptions establish the current conventions and physical constraints required for this analysis; they are not presented as a new topology or a comprehensive rederivation of LLC operating theory. The contribution lies in explicit transformations and quantitative descriptors that distinguish half-cycle correspondence, modal symmetry, reciprocal-gain mismatch, commutation feasibility, and circulating-energy cost. A physically consistent circuit description is therefore a prerequisite for the symmetry analysis, not a substitute for its validation. The present evidence is analytical and computational. Targeted hardware validation and an experimentally informed Pareto-based design methodology are planned as a subsequent stage, not results claimed in this manuscript.

Position Within the Resonant-Converter Research Line

Table 1 clarifies the progression from topology-specific LLC analysis and resonance-centered design to the present topology-specific symmetry formulation. The comparison is included to distinguish the new contribution from the authors' earlier design, optimization, quasi-resonant, and regime-level symmetry studies.

Table 1. Progression from earlier resonant-converter studies to the present LLC contribution.

Prior Contribution	Established Result	LLC-Specific Gap Addressed Here
Classical FHA and time-domain analysis [1–16]	Gain, impedance, peak-gain, and explicit operating-mode descriptions.	No common quantitative separation of reciprocal, half-cycle, modal, and commutation symmetry.
Resonant-converter symmetry studies [17–19]	Family- or regime-level symmetry and extended-duality language.	No event-aligned LLC state mapping or combined P/O/N-ZVS-circulation screen.
LLC design and efficiency optimization [23–26]	Wide-range design, tank-current inspection, peak-gain placement, and efficiency optimization.	Objectives are optimized, but symmetry breaking is not quantified as an independent diagnostic layer.
Quasi-boundary and QR-PWM studies [28,29]	Resonance-centered design and local ZVS/ZCS commutation symmetry.	Do not address LLC two-resonance reciprocal detuning and multimodal P/O/N structure.
Related LLC studies [20–22,27]	Operating-regime evaluation and model-based optimization.	No formal half-cycle operator, event-aligned residual, or modal-symmetry metric.
Present work	LLC-specific Z2 half-cycle mapping, reciprocal-detuning residual, P/O/N words, signed ZVS margin, circulation metric, convergence audit, and LTspice spot checks.	Hardware validation and calibrated magnetic/thermal loss models remain for future work.

2. LLC Converter and Analytical Models

2.1. Representative Converter and Normalized Variables

Figure 1 defines the circuit, terminal labels, and current conventions used throughout the analysis. Panel (a) provides fixed positive references. Panels (b1–b3) and (c1–c3)

enumerate the P/O/N secondary states for the two potentially enabled primary channels; their order is not a universal time sequence. Dashed i_r and i_m arrows denote signed references, whereas solid $|i_p|$ and $|i_s|$ arrows indicate the conduction direction for the stated secondary state. Panels (d_1–d_4) distinguish output-capacitance charge transfer from conditional body-diode clamping during dead time. The representative chronology and the model boundary are described in Sections 2.3.1 and 2.3.2.

The series and lower resonant frequencies are

$$f_r = \frac{1}{2\pi\sqrt{L_r C_r}}, f_m = \frac{1}{2\pi\sqrt{(L_r + L_m)C_r}} \quad (1)$$

Here, f_r is the series-resonant frequency of L_r and C_r, whereas f_m is the lower resonance associated with $L_r + L_m$ and C_r. The inductances L_r and L_m are the series-resonant and transformer magnetizing inductances, respectively (H); C_r is the series-resonant capacitance (F). Both frequencies are expressed in Hz, with kHz used in the numerical tables.

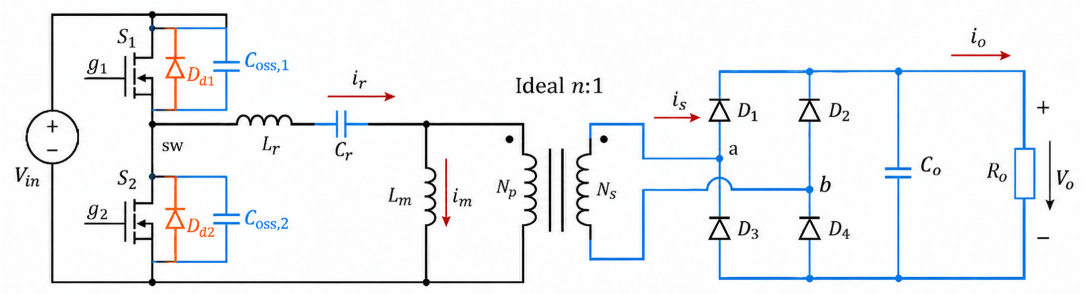

(**a**)

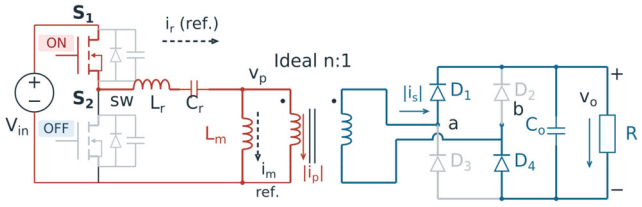

(**b1**)

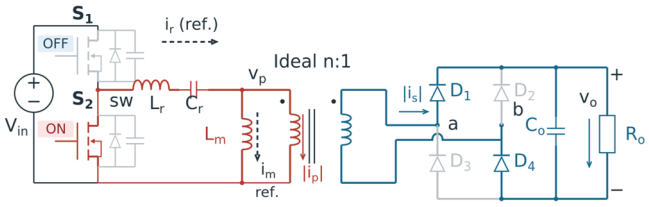

(**c1**)

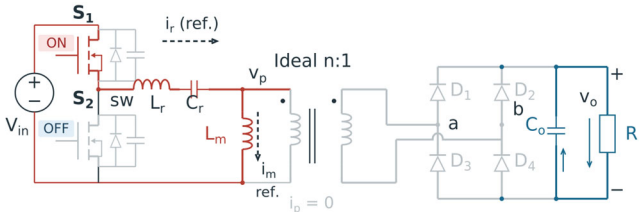

(**b2**)

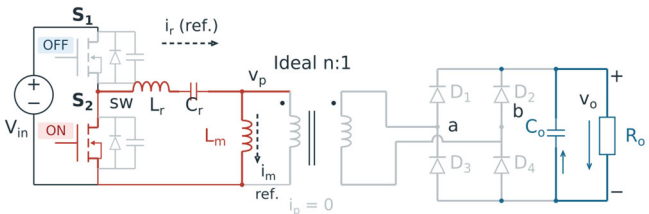

(**c2**)

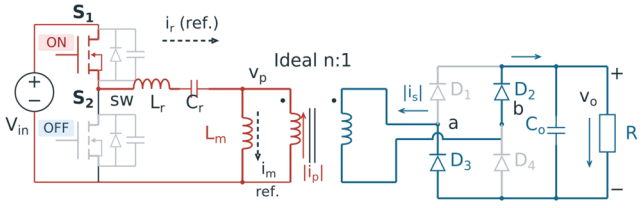

(**b3**)

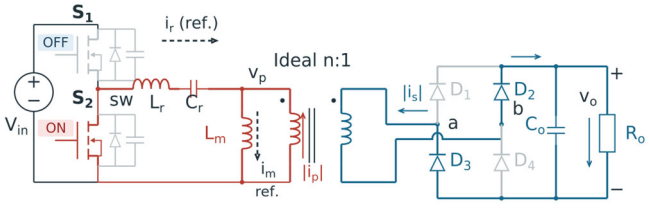

(**c3**)

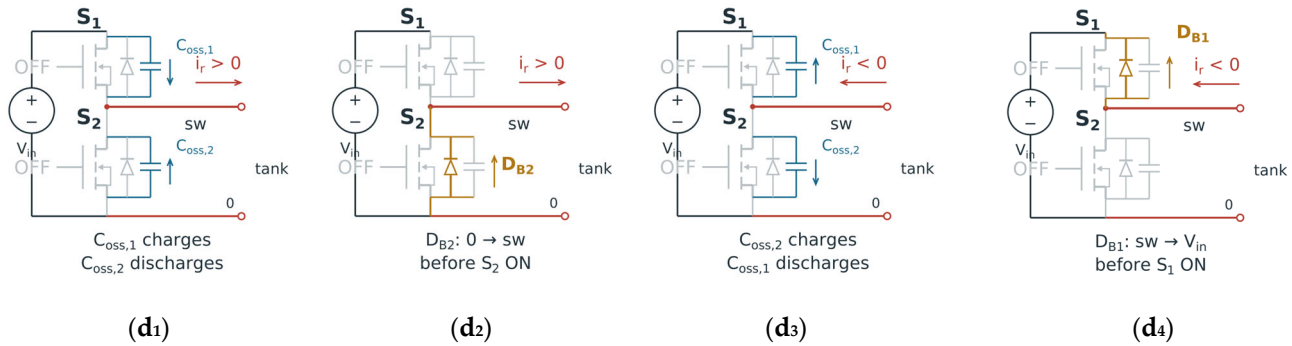

(**d**1) (**d**2) (**d**3) (**d**4)

Figure 1. Half-bridge LLC converter: reference circuit, secondary conduction states, and qualitative dead-time commutation . (**a**) Reference circuit with primary MOSFETs S_1 and S_2, gate commands g_1 and g_2, intrinsic body diodes D_{B1} and D_{B2}, and output capacitances $C_{oss,1}$ and $C_{oss,2}$. The magnetizing inductance L_m is in parallel with the ideal transformer primary. The fixed references satisfy $i_p = i_r - i_m$ and $i_s = ni_p$, where $n = N_p/N_s$. A negative current flows opposite to its reference arrow. The primary and secondary returns are galvanically isolated; the crossing of the secondary leads is not a junction. (**b**1–**b**3) S_1 ON and S_2 OFF; (**c**1–**c**3) S_2 ON and S_1 OFF. P denotes $i_p > 0$ with D_1/D4 conducting, N denotes $i_p < 0$ with D_2/D_3 conducting, and O denotes $i_p = i_s = 0$ with $i_r = i_m$ and Co supplying Ro. Dashed arrows are fixed positive references; solid current arrows show the stated conduction direction. Red identifies the enabled primary path, blue the rectified-output path, and grey an unhighlighted branch. ON/OFF refers to the channel command, not an exclusive division of current between channel and body diode. These are state combinations, not chronological steps. For an ideally half-cycle-symmetric solution, (**b**1) pairs with (**c**3), (**b**2) with (**c**2), and (**b**3) with (**c**1). (**d**1–**d**4) Qualitative dead-time commutation close-ups with both MOSFET channels OFF. For $i_r > 0$, (**d**1) shows $C_{oss,1}$ charging and $C_{oss,2}$ discharging during the downward switch-node transition; (**d**2) shows the conditional D_{B2} clamp before S_2 turn-on. For $i_r < 0$, (**d**3) shows the opposite capacitive transfer during the upward transition; (**d**4) shows the conditional D_{B1} clamp before S_1 turn-on. A clamp can occur only after the target rail is reached, and its interval may have zero duration. Body-diode conduction is not restricted to one side of series resonance. After gate turn-on, reverse current can use the enhanced channel; device-dependent channel/diode sharing is not resolved by these drawings. The O secondary state is not dead time. Equations (20)–(22) and Section 5.7 provide the separate commutation assessment; these sketches are not additional simulation results.

The normalized variables used throughout the paper are

$$f_n = \frac{f_s}{f_r}, m = \frac{L_m}{L_r}, Z_0 = \sqrt{\frac{L_r}{C_r}}, Q = \frac{Z_0}{R_{ac}} \quad (2)$$

Here, f_n is the switching frequency divided by the series-resonant frequency, $Q = \frac{1}{R_{ac}}\sqrt{\frac{L_r}{C_r}}$ is the load-dependent quality factor, and $m = L_m/L_r$ is the magnetizing-to-series inductance ratio. Q and m are dimensionless; they are not switching states.

The switching frequency f_s is in Hz and its period is $T_s = 1/f_s$ in seconds. The characteristic impedance Z_0 and equivalent load resistance R_{ac} are in ohms; consequently f_n, m and Q are dimensionless. The symbol Q here is a load-dependent quality factor, not an electrical charge.

Here, R_{ac} is the load resistance referred to the resonant-tank fundamental. With the common full-wave rectifier approximation,

$$R_{ac} = \frac{8n^2}{\pi^2} R_o \quad (3)$$

In Equation (3), R_o is the physical DC load resistance (Ω), and R_{ac} is its primary-referred fundamental-frequency equivalent (Ω). The dimensionless turns ratio is $n = N_P/N_s$,

where N_p and N_s are the primary and secondary winding turns. Thus, n is a transformer ratio, not the normalized frequency f_n.

The symmetry model uses centered variables, specifically, $v_g = v_{sw} - V_{in}/2$ and $v_{Cr} = v_{Cr,physical} - V_{in}/2$, for the balanced half bridge. Thus, the odd half-cycle mapping applies to the centered capacitor voltage, not to a physical capacitor voltage that includes its DC bias. The supplementary terminal-filter model retains the physical capacitor voltage directly.

All voltages in this convention are in volts: V_{in} is the imposed DC input voltage; v_{sw} is the instantaneous half-bridge switch-node voltage relative to the primary return; and $v_{Cr,physical}$ is the actual resonant-capacitor voltage. The centered variables v_g and v_{Cr} subtract the stated DC bias. They must not be interchanged with the physical terminal voltages when interpreting the symmetry transformation.

The parameter set in Table 2 is illustrative rather than tied to a particular commercial design. It is selected to create a nominal 400 V to 48 V isolated conversion example with $f_r = 100$ kHz and a moderate magnetizing ratio $m = 5$.

Table 2. Representative converter parameters and the corresponding model detail used in the low-order and independent LTspice calculations.

Parameter or Model Element	Symbol	Low-Order Analytical/Switched Model	Frozen LTspice Implementation/Explicitly Separate Extension
Input voltage	V_{in}	400 V	400 V
Series resonant frequency	f_r	100 kHz	100.000 kHz from L_r and C_r
Series resonant inductance	L_r	60 µH, ideal energy storage	60 µH with 45 mΩ series resistance
Series resonant capacitance	C_r	42.217 nF, ideal	42.217 nF with 35 mΩ ESR
Magnetizing inductance	L_m	300 µH	300 µH coupled-transformer magnetizing branch
Magnetizing ratio	$m = L_m/L_r$	5	5
Transformer ratio	$n = N_p/N_s$	4.1667, ideal	4.1667; k = 0.9995; R_p/R_s = 80/12 mΩ
Output capacitance	C_o	470 µF	470 µF with 25 mΩ ESR
Load resistance at Q = 0.4	R_o	6.6973 Ω	6.6973 Ω
Primary switch conduction	S_1, S_2	Ideal switches	45 mΩ channel resistance and body-diode paths
MOSFET output capacitance	C_{oss}	200 pF/device constant charge screen	Nonlinear $Q(V)$ at f_n = 0.80; 278 pF/device charge-equivalent at f_n = 0.45 and 1.25
Dead time and gate edge	t_d	200 ns charge screen	200 ns dead time; 20 ns rise/fall time
Rectifier	D_1–D_4	Ideal clamp with controlled drop mismatch cases	Dynamic diode model: R_s = 20 mΩ, C_{j0} = 40 pF, T_t = 5 ns
Integration software	-	Python 3/NumPy/SciPy; fixed-step RK4 plus adaptive event refinement	LTspice 24.1.7 for Windows, Gear integration
Lower resonant frequency	f_m	40.825 kHz	40.825 kHz
Characteristic impedance	Z_0	37.699 Ω	37.699 Ω
Practical MOSFET reference	S_1, S_2	Ideal switches remain in the symmetry solver	IPP65R190CFD7A [34]; separate datasheet-informed calculations in Sections 5.8 and 5.9. Not the device used in the archived generic LTspice netlists.
Practical rectifier reference	D_1–D_4	Ideal clamp remains in the symmetry solver	Four STTH30R02DJF diodes [35]; separate conduction and recovery-sensitivity calculations.

Added input filter (Section 5.9 only)	L_f, C_f	Absent in the frozen model	20 µH/10 µF; inductor resistance 50 mΩ and capacitor ESR 100 mΩ. Illustrative, not optimized.

The named devices in the last rows of Table 2 define an additional practical reference, not a retrospective identification of the generic LTspice elements. Their parameters and the assumptions used to estimate losses are considered separately in Sections 5.8 and 5.9. The analytical symmetry maps and the archived LTspice comparison are retained unchanged.

The two resonance landmarks partition the normalized-frequency plane into three broad screening regions. For $f_n < f_m/f_r = 0.408$, the input impedance may become capacitive and unfavorable for ZVS. For $f_m/f_r < f_n < 1$, boost-like operation contains several P/O/N mode families. For $f_n > 1$, the gain is normally below unity and the tank is commonly inductive. These are not universal hard boundaries: the input phase, load factor Q, rectifier state, parasitics, and exact modal sequence remain operating-point-dependent.

2.2. FHA Gain and Reciprocal-Detuning Decomposition

Under the fundamental harmonic approximation, the normalized voltage gain used in the present study is

$$M_{\mathrm{FHA}}(f_n, Q, m) = \left\{ \left[1 + \frac{1}{m}\left(1 - \frac{1}{f_n^2}\right) \right]^2 + Q^2 \left(f_n - \frac{1}{f_n} \right)^2 \right\}^{-1/2} \tag{4}$$

M_{FHA} is the dimensionless fundamental-harmonic approximation to the normalized DC conversion gain. The normalization is $M = 2nV_o/V_{in}$ for the half bridge, where V_o denotes the period-mean output voltage. The time-domain quantity M_{TD} in the results uses the same normalization, with the mean taken from the switched-model output. The arguments f_n, Q and m retain the definitions in Equation (2).

The gain satisfies the expected LLC dependence on normalized frequency, load, and magnetizing ratio. For finite m, the lower resonance creates the familiar gain peak at light and medium load. At family level, LLC occupies a bridge position between series-dominant transfer and output-parallel shaping; at operating level, that bridge is resolved by the five topology-specific symmetry coordinates and their quantitative metrics. Figure 2 summarizes this hierarchy.

A resonance-centered reflection is most naturally written in the logarithmic coordinate

$$\xi_f = \ln f_n, D_F : \xi_f \mapsto -\xi_f, \text{ or equivalently } f_n \mapsto f_n^{-1} \tag{5}$$

The dimensionless coordinate ξ_f is the natural logarithm of f_n; D_F denotes the frequency-reflection operation. Its image has normalized frequency $1/f_n$, with Q and m held fixed. The superscript −1 denotes a reciprocal, not the negative-secondary state N.

The squared denominator of (4) can be separated as

$$D(f_n) = A_m^2(f_n) + B_Q^2(f_n) \tag{6}$$

where

$$A_m(f_n) = 1 + \frac{1}{m}(1 - f_n^{-2}), B_Q(f_n) = Q(f_n - f_n^{-1}) \tag{7}$$

The function D is the squared denominator of the FHA gain, so that $M_{FHA} = 1/\sqrt{D}$. The dimensionless functions A_m and B_Q are, respectively, the finite-magnetizing contribution and the load-weighted reactive-detuning contribution to that denominator. Their subscripts identify the parameters on which they depend. The function D is distinct from the duty fraction D_k introduced for event alignment.

https://doi.org/10.3390/sym18091561

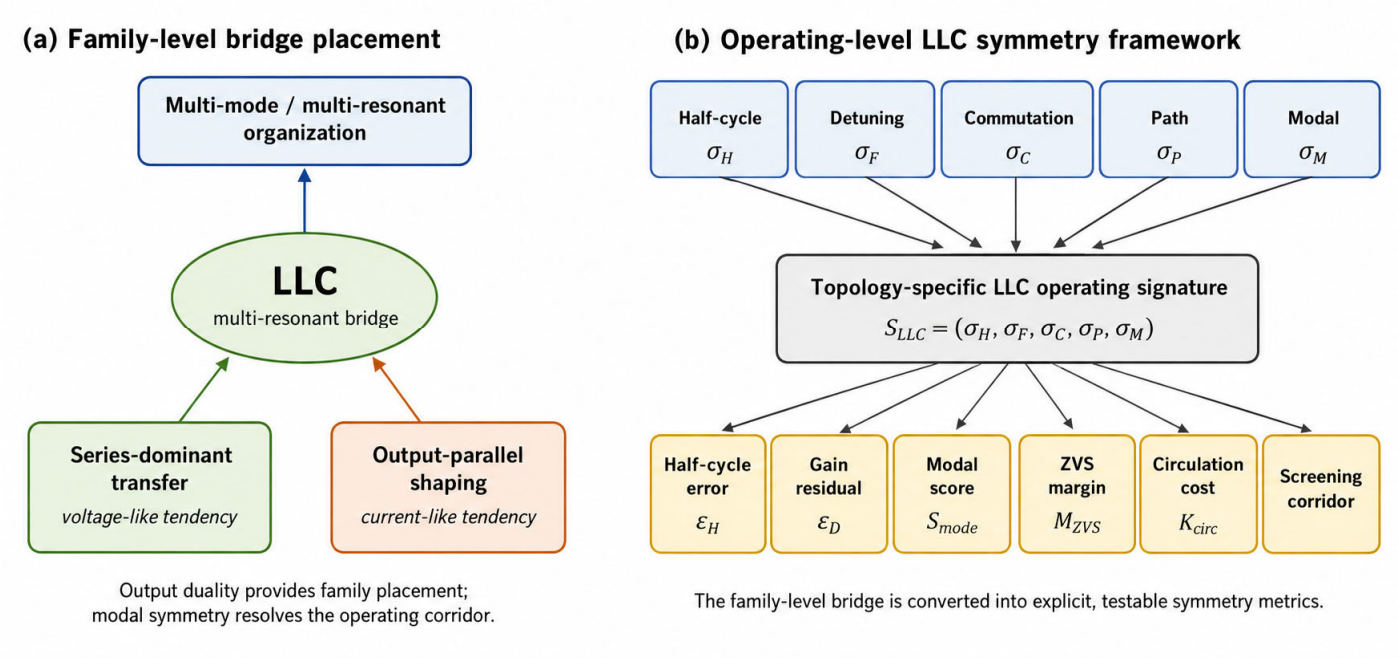

Figure 2. Hierarchical interpretation of LLC symmetry. (**a**) Family-level placement of LLC as a multi-resonant bridge between series-dominant transfer and output-parallel shaping. (**b**) Operating-level specialization into half-cycle, detuning, commutation, path, and modal descriptors with quantitative screening metrics.

The reactive term is invariant in magnitude under reciprocal reflection,

$$B_Q^2(f_n) = B_Q^2(f_n^{-1}) \tag{8}$$

whereas the finite-magnetizing term is not. The exact squared-denominator difference is

$$\Delta_m(f_n) = D(f_n) - D(f_n^{-1}) = \frac{1}{m}(f_n^2 - f_n^{-2})\left[2 + \frac{1}{m}(2 - f_n^2 - f_n^{-2})\right] \tag{9}$$

Here, Δ_m is the signed difference between the squared gain denominators at reciprocal normalized frequencies. It is dimensionless and is not a voltage difference or a manufacturing tolerance. The comparison retains the same m and Q at both frequencies.

Equation (9) isolates finite magnetizing inductance as a structural symmetry-breaking mechanism. In the series-resonant limit $m \to \infty$, $\Delta_m \to 0$ and the reciprocal gain symmetry is restored. A dimensionless residual used below is

$$\varepsilon_D(f_n) = \frac{|M_{\mathrm{FHA}}(f_n) - M_{\mathrm{FHA}}(f_n^{-1})|}{1/2\,[M_{\mathrm{FHA}}(f_n) + M_{\mathrm{FHA}}(f_n^{-1})]} \tag{10}$$

The ε_D is the dimensionless reciprocal-gain residual. The numerator is the absolute difference of the two FHA gains, and the denominator is their arithmetic mean; the omitted arguments Q and m are identical in both gains. A percentage is obtained by multiplying ε_D by 100. This measures gain correspondence, not half-cycle waveform mismatch.

2.3. Piecewise P/O/N Switched Model

The FHA model is complemented by a topology-level switched model. In Figure 1, $v_{ab} = v_a - v_b = v_p/n$, and $i_s = n(i_r - i_m)$ is positive from terminal a into the rectifier. These fixed physical references apply for either enabled primary switch. Within each half-cycle, the rectifier-referred transformer voltage is assigned from three physically distinct stages:

- P: positive-secondary conduction ($i_s > 0$); $v_p = n(v_o + V_{D,p})$, and the load path is a → D_1 → (C_o ‖ R_o) → D_4 → b. D_1 and D_4 conduct, while D_2 and D_3 are blocked. $V_{D,p}$ denotes the total drop of the conducting diode pair.

- O: secondary open; all four output diodes are blocked and the ideal-transformer load current is zero, so $i_r = i_m$. The primary voltage is generally nonzero: $v_p = L_m\, \mathrm{d}(i_m)/\mathrm{d}t = [L_m/(L_r + L_m)](v_g - v_{Cr})$ in the lossless centered model. C_o supplies the load during this interval.
- N: negative-secondary conduction ($i_s < 0$); $v_p = -n(v_o + V_{D,n})$, and the load path is b → D_2 → (C_o ∥ R_o) → D_3 → a. D_2 and D_3 conduct, while D_1 and D_4 are blocked. $V_{D,n}$ likewise denotes the total pair drop.

Current signs and diode selection: The positive reference of i_r is from the switching node sw toward L_r for both gate states. Therefore $i_r < 0$ means actual current toward the switching node sw, without changing the reference definition. P requires $i_r > i_m$, O has $i_r = i_m$ with the secondary open, and N requires $i_r < i_m$. These comparisons involve signed currents: N does not require $0 < i_r < i_m$, which is only one possible case. Enabling S_2 does not impose an instantaneous reversal of i_r or an automatic exchange of the rectifier pair. Within an established channel-conduction interval, the drain-to-source current is i_r for S_1 and $-i_r$ for S_2; reverse channel conduction is distinct from body-diode conduction.

Half-cycle pairing is a different comparison. Under the ideal half-cycle-symmetric periodic solution discussed in Section 3.2, i_r, i_m, i_p and i_s at $t + T_s/2$ are the negatives of their values at t. The physical secondary state consequently exchanges P↔N while O↔O, so Figure 1(b_1) pairs with (c_3), (b_2) with (c_2), and (b_3) with (c_1). The diode identities and terminal labels remain fixed. This symmetry correspondence does not assert that all six combinations occupy a finite interval at every operating point, nor that current reverses at the gate–command edge.

For P or N conduction, with $s \in \{+1, -1\}$,

$$\begin{aligned} \frac{di_r}{dt} &= \frac{v_g - v_{Cr} - sn\left(v_o + V_{D,s}\right)}{L_r}, \\ \frac{di_m}{dt} &= \frac{sn\left(v_o + V_{D,s}\right)}{L_m}, \\ \frac{dv_{Cr}}{dt} &= \frac{i_r}{C_r}, \\ \frac{dv_o}{dt} &= \frac{n|i_r - i_m| - v_o/R_o}{C_o}, s \in \{-1, +1\}. \end{aligned} \tag{11}$$

In Equation (11), t is physical time (s); i_r and i_m are the resonant-branch and magnetizing currents (A); v_g, v_{Cr} and v_o are the centered bridge excitation, centered resonant-capacitor voltage and instantaneous output voltage (V). The output capacitance C_o is in farads and R_o is in ohms. The symbol s selects the secondary clamp: +1 for P and −1 for N, independently of which primary switch is enabled. The pair drop $V_{D,s}$ equals $V_{D,p}$ in P and $V_{D,n}$ in N (V). The term $n|i_r - i_m|$ is the rectified secondary current delivered to the output, while v_o/R_o is the load current. Their difference is the output-capacitor charging current.

During the O stage, the secondary current is zero and $i_r = i_m$:

$$\begin{aligned} \frac{di}{dt} &= \frac{v_g - v_{Cr}}{L_r + L_m}, \\ \frac{dv_{Cr}}{dt} &= \frac{i}{C_r}, \\ \frac{dv_o}{dt} &= \frac{-v_o}{R_o C_o}. \end{aligned} \tag{12}$$

Here, i is the common current $i_r = i_m$ during O, not the secondary current, which is zero. The derivative d/dt is taken with respect to physical time; the current and voltage derivatives therefore have units A/s and V/s. The final equation describes discharge of C_o into R_o while the secondary is open. The same centered-voltage convention is used in both state systems.

Stage transitions are detected from the transformer-voltage clamp condition and the zero crossing of $i_r - i_m$. The model includes output-capacitor dynamics and therefore produces a self-consistent DC output. Nonlinear device capacitance C_{oss}, transformer coupling loss, winding resistance, and rectifier dynamics are omitted from this low-order solver but are introduced in the independent LTspice spot checks described in Section 4.2.

2.3.1. Chronology of the Representative Operating Families

Figure 1 enumerates conduction states; their time order is operating-point-dependent. The corresponding low-order waveforms for the three representative cases are presented later in Section 5.2. The following sequences use the fixed physical P/O/N convention before canonicalizing the negative half-cycle; a gate–command change does not force an instantaneous secondary-current reversal.

At $f_n = 0.45$, the positive-excitation sequence is P → O → N. The D_1/D_4 current reaches zero, the secondary opens with $i_r = i_m$, and D_2/D_3 subsequently establish the opposite clamp. The complementary physical sequence is N → O → P. At $f_n = 0.80$, the sequences are P → O and N → O. The O interval supplies no transformer load current; C_o supplies R_o. These sequences correspond to the representative low-order waveforms presented in Section 5.2 and do not imply completed ZVS.

At $f_n = 1.25$, negative-secondary conduction initially persists into positive excitation, followed by positive-secondary conduction: the reported compressed words are NP and PN. A resolved short open interval gives NOP and PON instead. Section 3.3 distinguishes this NP-family classification from the full extracted word and states the extraction threshold; the corresponding low-order and circuit-level results presented in Section 5 retain the reported classification. In all three cases, current-zero and transformer-clamp events determine the secondary sequence, while commutation is assessed separately.

2.3.2. Physical Commutation and the Model Boundary

After the outgoing channel is turned OFF, both channels are OFF during dead time. For $i_r > 0$, the downward transition charges $C_{oss,1}$ and discharges $C_{oss,2}$; D_{B2} can clamp the lower rail only after that rail is reached. For $i_r < 0$, the upward transition charges $C_{oss,2}$ and discharges $C_{oss,1}$; D_{B1} can then clamp the upper rail. The incoming channel may subsequently carry reverse current before the current crosses zero. The diode-clamp interval is conditional and may have zero duration. The ON-state drawings identify an available channel path; they do not resolve device-dependent channel/body-diode current sharing or imply that an unhighlighted body diode has zero current in a real device. Figure 1(d_1–d_4) shows the channels-OFF commutation stages, whereas Equations (20)–(22) and the circuit-level results presented in Section 5.7 provide the separate signed-charge screen and validation check. No dead-time state is added to the low-order solver by this explanation.

3. Symmetry and Extended-Duality Framework

3.1. LLC Operating Signature

Two descriptive levels are deliberately separated. At family level, LLC is assigned a hybrid, region-dependent output character because series-dominant transfer and magnetizing-branch shaping coexist. That family label is structural rather than state-dependent. The operating-level signature below therefore retains the five coordinates that change with frequency, commutation, feasible paths, half-cycle mapping, and modal sequence. This hierarchy prevents the broad output-duality classification from being conflated with the quantitative mode-by-mode symmetry metrics.

The previously proposed multidimensional resonant-converter signature is specialized to the LLC topology as

$$s_{\mathrm{LLC}} = (\sigma_H, \sigma_F, \sigma_C, \sigma_P, \sigma_M) \tag{13}$$

The s_{LLC} is a five-component operating signature, not the scalar clamp sign s of Equation (11). Its categorical coordinates σ_H, σ_F, σ_C, σ_P and σ_M describe half-cycle parity, detuning side, commutation tendency, path feasibility and modal sequence, respectively. These labels have no physical units. In Table 3, sgn denotes the sign function, which distinguishes negative, zero and positive logarithmic detuning.

The descriptors are summarized in Table 3. The source-domain coordinate of the general framework is replaced by half-cycle parity because the present topology has a fixed voltage-fed half bridge. Continuous metrics accompany the categorical descriptors; the signature is therefore a compact label, not a substitute for the waveform solution.

Table 3. Topology-specific LLC symmetry descriptors.

Descriptor	Meaning	Representative Quantitative Evidence
σ_H	Positive/negative half-cycle parity	ε_H
σ_F	Sub-resonant, resonant, or super-resonant side	$\mathrm{sgn}(\ln f_n)$ and ε_D
σ_C	Commutation tendency	m_{ZVS}
σ_P	Physically valid resonant/rectifier path	P/O/N stage feasibility
σ_M	Ordered conduction-mode structure	W and S_M

3.2. Half-Cycle Z_2 Symmetry

Let the state column vector contain, in this order, the resonant current, magnetizing current, centered resonant-capacitor voltage, and output voltage:

$$x(t) = [i_r(t) \quad i_m(t) \quad v_{Cr}(t) \quad v_o(t)]^T \tag{14}$$

For an ideal balanced bridge, identical rectifier drops, and periodic steady state, the half-cycle transformation is

$$H: x\left(t + \frac{T_s}{2}\right) = S_H x(t), S_H = \mathrm{diag}(-1, -1, -1, +1) \tag{15}$$

The state vector x contains the four instantaneous quantities in the order shown in Equation (14); its superscript T denotes transpose, not a period or temperature. The operator H performs the half-cycle transformation, and S_H is its dimensionless diagonal sign matrix; diag lists the matrix diagonal. The time shift is $T_s/2$. In the identity stated below, I is the identity operation and Z_2 denotes the two-element symmetry in which applying the transformation twice returns the original state.

At the same time, $v_g(t + T_s/2) = -v_g(t)$ and the conduction labels are exchanged through

$$\rho(P) = \mathrm{N}, \rho(N) = \mathrm{P}, \rho(O) = \mathrm{O} \tag{16}$$

The symbol ρ is a permutation of the stage labels: it exchanges P and N and leaves O unchanged. It has no physical unit and does not alter the fixed diode numbering in Figure 1.

Because $H^2 = I$, the ideal half-cycle mapping forms a Z_2 symmetry. The output voltage is even under the mapping, while the resonant and magnetizing currents and the resonant-capacitor voltage are odd.

To remove artificial mismatch caused by splitting an imbalanced waveform at a nominal $T_s/2$, the revised metric is aligned to the actual commanded switching event. Let t_k denote the beginning of cycle k. The commanded positive-to-negative transition divides that cycle into intervals of duration $D_k T_s$ and $(1 - D_k)T_s$. The two intervals are resampled independently on $\tau \in [0,1]$, producing $x_k^+(\tau)$ and $x_k^-(\tau)$. The comparison error is $e_h(\tau) = x_k^-(\tau) - S_h x_k^+(\tau)$, where $S_h = \mathrm{diag}(-1,-1,-1,+1)$.

The dimensionless half-cycle residual is

$$\varepsilon_{\mathrm{H}} = \sqrt{\frac{\int_0^1 e_{\mathrm{H}}{}^{\mathrm{T}} \mathrm{W} e_{\mathrm{H}} \mathrm{d}\tau}{\int_0^1 (x^+)^{\mathrm{T}} \mathrm{W} x^+ \mathrm{d}\tau}} \tag{17}$$

In Equation (17), x^+ and x^- denote the two resampled state trajectories, with the cycle index k suppressed, and e_{H} is their state-by-state difference after applying S_{H}. The variable τ is dimensionless local time from 0 to 1 within each resampled interval. The earlier t_{k} is the start of cycle k (s), and D_{k} is its dimensionless positive-command duration divided by T_{s}. These cycle indices are local to the alignment construction.

The diagonal matrix W scales each state by its full-period RMS value, so ε_{h} is dimensionless and no single high-amplitude state dominates. Here, $\varepsilon_{\mathrm{h}} = 0$ denotes exact numerical half-cycle correspondence at the adopted model level; it does not imply reciprocal-gain equality, low circulating energy, or ZVS. The revised values are obtained after fixed-step warm-up by adaptive integration with explicit P/O/N event location.

$$W = \mathrm{diag}\left(I_{r,\mathrm{rms}}^{-2}, I_{m,\mathrm{rms}}^{-2}, V_{Cr,\mathrm{rms}}^{-2}, V_{o,\mathrm{rms}}^{-2}\right)$$

The quantities $I_{\mathrm{r,rms}}$, $I_{\mathrm{m,rms}}$, $V_{\mathrm{Cr,rms}}$ and $V_{\mathrm{o,rms}}$ are the full-period root-mean-square values of the corresponding states, in A, A, V and V. Their inverse squares define the diagonal weights in W, making the quadratic products in Equation (17) dimensionless. The subscript rms means root mean square; the capitalization RMS used elsewhere denotes the same operation. The transpose in each quadratic product is an algebraic operation, not another time shift.

3.3. Modal Words and Modal Symmetry

The ordered sequence of P/O/N intervals in the positive half-cycle is represented by a modal word

$$W_P = M_1 M_2 \cdots M_k, M_j \in \{P, O, N\} \tag{18}$$

W_{P} is the ordered modal word of the positive half-cycle and W_{N} denotes the corresponding negative-half-cycle word. Each M_{j} is one retained P, O or N stage; j is its position and k is the number of stages in the word. These are discrete labels, not gain values. A modal word, denoted generically by W in Table 3, is distinct from the weighting matrix W in Equation (17).

The negative-half-cycle word is converted to the positive canonical alphabet with (16), giving $\rho(W_N)$. A normalized edit-distance score is

$$S_M = 1 - \frac{d_L(W_P, \rho(W_N))}{\max(L_P, L_N)}, L_P = \mathrm{len}(W_P), L_N = \mathrm{len}(W_N) \tag{19}$$

where d_L is the Levenshtein distance. $S_M = 1$ means that the two half-cycles share the same canonical modal sequence even if their continuous waveforms are not exactly identical. This distinction is useful: modal symmetry is discrete, whereas ε_H captures continuous amplitude and timing mismatch.

L_{P} and L_{N} are the numbers of letters in W_{P} and W_{N}; len counts those letters and max selects the greater word length. They are not inductances. The Levenshtein distance d_{L} counts the minimum single-letter insertions, deletions and substitutions needed to match the words. The dimensionless score S_{M} is the modal-similarity quantity labeled S_{mode} in the numerical tables. Canonicalization applies ρ to every letter of W_{N} before comparison.

The P↔N exchange in Equation (16) is used only to canonicalize the negative-half-cycle word. Figure 1 retains fixed physical a/b polarity and diode identities: a canonical P in the reflected negative half-cycle represents physical N conduction through D_2 and D_3. The highlighted paths describe ideal conduction intervals, not brief parasitic recovery or displacement currents.

The mode label NP is used for the reduced-model family and for the archived circuit extraction at its stated 0.05 A secondary-current threshold. NOP denotes a variant with a resolved short O interval between N and P, not an identical discrete word. Both archived extracted words are reported as NP; the NP-family terminology accommodates the short-interval variant without replacing the archived classification. In the negative half-cycle, the physical P/N labels are first exchanged according to Equation (16).

3.4. Commutation Symmetry and the Cost of Symmetry

A transition-oriented ZVS charge margin is defined at each commanded bridge transition k. The orientation factor s_k specifies the current reference for the required node movement; assistance is determined by the sign of the product of s_k and the resonant current, not by s_k alone. Retaining this sign prevents a large current of the wrong polarity from being counted as available commutation charge.

$$Q_{av,k} = s_k\, i_r(t_k^-)\, t_d, \quad s_k \in \{-1, +1\} \tag{20}$$

$Q_{av,k}$ is the signed available commutation charge (C), t_d is the dead time (s), and $i_r(t_k^-)$ is the resonant current immediately before transition k (A). Here k indexes a bridge transition rather than a complete cycle. For the current reference in Figure 1, s_k is −1 for the upward node transition and +1 for the downward transition. Thus, an assisting current gives positive available charge; a current of the wrong polarity is not converted to an absolute magnitude.

For voltage-dependent MOSFET capacitance, the required charge is the sum of the two device output-charge integrals over the commutated voltage. The constant-capacitance screening model is recovered as the final approximation in Equation (21).

$$Q_{\mathrm{req,k}} = \int_0^{V_{\mathrm{in}}} C_{\mathrm{oss,H}}(v)\, dv + \int_0^{V_{\mathrm{in}}} C_{\mathrm{oss,L}}(v)\, dv \approx 2\, C_{\mathrm{oss}}\, V_{\mathrm{in}} \tag{21}$$

$Q_{\mathrm{req,k}}$ is the required node-transition charge (C). The functions $C_{\mathrm{oss,H}}(v)$ and $C_{\mathrm{oss,L}}(v)$ are the voltage-dependent output capacitances of the high-side and low-side MOSFETs (F); v is the integration voltage (V). In the constant-capacitance approximation, C_{oss} is the value for one device and the factor two accounts for both devices. These charge quantities are distinct from the dimensionless load factor Q.

The normalized transition margin and its worst-transition value are

$$m_{\mathrm{ZVS}} = \min_{\mathrm{k}}\left(\frac{Q_{\mathrm{av,k}} - Q_{\mathrm{req,k}}}{Q_{\mathrm{req,k}}}\right) \tag{22}$$

The minimum is taken over the two transitions. A positive value indicates sufficient charge and correct polarity under the adopted charge model. This is a screening condition, not device-level proof; the LTspice spot checks additionally measure residual drain-source voltage immediately before turn-on. Nonlinear output-charge and dead-time effects in LLC converters require this transition-specific treatment [36].

Here, m_{ZVS} is the dimensionless worst-transition relative charge margin, with min selecting the smaller value over the two bridge transitions. Zero marks equality between available and required charge; a negative value fails the adopted screen. The subscript ZVS distinguishes this margin from the inductance ratio m. The drain–source voltage V_{DS} used in the separate circuit check is in volts.

Soft-switching symmetry is not free. The mean stored reactive energy is

$$\overline{E}_r = \left\langle \frac{L_r i_r^2}{2} + \frac{L_m i_m^2}{2} + \frac{C_r v_{Cr}^2}{2} \right\rangle_T \tag{23}$$

$\overline{E}_r$ is the period-mean reactive-energy indicator (J). Angle brackets with subscript T denote averaging over one complete switching period, with T = T_s here. The three terms represent the two inductive stores and the resonant-capacitor contribution evaluated with the centered v_{Cr} state. The overbar denotes this period average. The capacitor term uses

the centered voltage rather than its physical biased value; $\overline{E}_r$ is not a sum of all physical filter and device energies.

And the normalized circulation cost is

$$\kappa_{\mathrm{circ}} = \frac{\omega_s \overline{E}_r}{P_o} \tag{24}$$

Here, $\omega_s = 2\pi f_s$ is the switching angular frequency (rad/s), and P_o is the average active output power (W), also denoted P_{out} in the results and loss budget. For the resistive load it is the period average of v_o^2/R_o, not generally the square of the mean voltage divided by R_o. The circulation index κ_{circ} is dimensionless; it is neither an efficiency nor a loss fraction.

The metric estimates how much reactive energy is cycled per delivered active power. It is intended for relative screening across operating points with the same topology and parameter basis.

4. Numerical Methodology

4.1. Frozen Low-Order and Event-Refined Workflow

The frozen workflow combines analytical FHA mapping, a piecewise P/O/N switched solver, adaptive event refinement, and hard validation gates. The analytical FHA and reciprocal-residual curves cover $0.35 \le f_n \le 1.65$, whereas the switched-model/modal/corridor grid covers $0.35 \le f_n \le 1.55$ and $Q = 0.2$–0.9. Grid, asymmetry, and tolerance runs use 720 cycles and 300 fixed fourth-order Runge–Kutta (RK4) steps per cycle; selected-point waveforms use 1300 cycles and 720 steps per cycle. Headline ε_H values are recomputed after the fixed-step warm-up using 20 adaptive polishing cycles and 1600 interpolation points with explicit event location.

For each operating point, the workflow exports time-domain and FHA gain, canonical positive and negative modal words, event-aligned ε_H, periodic-steady-state residual ε_{PSS}, transition-specific oriented currents and ZVS margins, RMS resonant and magnetizing currents, RMS resonant-capacitor voltage, circulation index κ_{circ}, and reciprocal-frequency residual ε_D.

The periodic-steady-state residual ε_{PSS} compares successive complete switching cycles, whereas ε_H compares mapped intervals within a cycle and ε_D compares gains at reciprocal frequencies. All three are dimensionless but test different kinds of correspondence; their subscripts must therefore be retained.

Four controlled symmetry-breaking cases are applied at $Q = 0.4$ over $0.55 \le f_n \le 1.45$: (A) a +2% positive/negative bridge-amplitude imbalance, (B) a +0.5% duty imbalance, (C) a 0.5 V negative-rectifier-drop mismatch, and (ABC) the simultaneous combination of cases A, B, and C. Central-range summaries use $0.65 \le f_n \le 1.35$. Each case changes only the stated mechanism; the experiment is a sensitivity diagnostic, not a complete manufacturing or thermal tolerance model.

Three representative points are retained for detailed interpretation: (f_n, Q) = (0.45, 0.40), (0.80, 0.40), and (1.25, 0.40). They span near-lower-resonance high-gain PON operation, an intermediate PO regime, and super-resonant NP operation.

A 300-run deterministic-seed tolerance study is performed at fixed f_s = 80 kHz. The primitive physical quantities L_r and C_r vary independently by ±5%, L_m and R_o by ±10%, C_{oss} by ±20%, and dead time by ±10%, all with uniform bounded distributions. The actual f_r, f_n, and Q are derived for every draw; Q is not sampled independently of L_r and C_r. These independent uniform bounds are an exploratory local model and do not represent correlated manufacturing lots, temperature drift, or long-term ageing.

Numerical acceptance is based on separate audits rather than one mixed convergence test. A selected point is accepted when the normalized mismatch between the final two cycles satisfies $\varepsilon_{PSS} < 1 \times 10^{-2}$, the gain changes by less than 0.05% between the 1300- and

1800-cycle runs, the modal family is unchanged, and the event-aligned ε_H varies by less than 1×10^{-4} across 600–1800 interpolation points at 20 polishing cycles. The validation script fails if required exports are missing, ideal ε_H exceeds 0.005, tolerance-run ZVS becomes non-positive, or canonical modal pairing is lost.

4.2. Independent Parasitic-Aware LTspice Spot Checks

Independent circuit-level spot checks were executed in LTspice 24.1.7 for Windows at $f_n = 0.45$, 0.80, and 1.25 using the same 400 V, $Q = 0.4$, $m = 5$ converter. The netlists include complementary S_1/S_2 commands, 200 ns dead time, finite switch resistance, body-diode paths, resonant and winding resistances, capacitor ESR, transformer coupling $k = 0.9995$, and dynamic rectifier diodes. The $f_n = 0.80$ case uses a nonlinear MOSFET output-charge law; the $f_n = 0.45$ and 1.25 convergence-safe checks use its 400 V charge-equivalent value of 278 pF/device. Gear integration, precharged steady-state variables, and final-window averaging are used only to obtain stable circuit simulations and do not share code with the P/O/N solver.

The circuit-level ZVS indicator is residual V_{DS} immediately before the commanded turn-on; $V_{DS} < 0.05V_{in} = 20$ V is used as a transparent practical threshold. Electrical efficiency is computed from averaged input and output powers. It includes the explicit electrical losses in Table 2 but excludes gate-drive power, magnetic-core loss, saturation, thermal dependence, layout parasitics not represented by the netlist, and hardware measurement uncertainty. The LTspice results are therefore independent numerical validation, not experimental.

4.3. Separate Semiconductor-Loss and Terminal-Filter Extension

A separate lossy P/O/N model is used for the practical reference and the terminal-filter comparison. It retains the same L_r, C_r, L_m, transformer ratio, load, and nominal input voltage, but includes the selected MOSFET channel resistance and a datasheet-based piecewise-linear rectifier drop. The input filter and output capacitor are coupled states; filtering is not applied digitally to previously generated waveforms. The complete state equations and parameter provenance are given in Supplementary Material S2.

A 3000-cycle fixed-step warm-up provides an initial state. Periodic shooting and event-located DOP853 integration then establish the periodic solution, and are followed by 60 verification cycles. The final 16 cycles are sampled at 2048 points per cycle. Period averages are computed using 32-point Gaussian quadrature on smooth event intervals, avoiding numerical bias at bridge-current jumps. A separate check uses 120 verification cycles, twice the output sampling density, and a maximum step reduced from $T_s/100$ to $T_s/200$. These checks concern this extension only; they do not replace the convergence evidence for the original symmetry metric.

The extension resolves conduction and filter energy storage, but not dynamic C_{oss}, gate trajectories, reverse recovery, magnetic-core loss, or temperature feedback. Switching and gate-drive losses are therefore estimated separately from datasheet parameters under stated scenarios. The new numbers are neither manufacturer-macromodel LTspice results nor laboratory measurements. The original independent LTspice checks remain the circuit-level evidence for the earlier generic model.

5. Results

5.1. Gain Surface and Reciprocal-Frequency Duality

Figure 3 shows the FHA gain curves and the two resonant landmarks. The lower resonance is visible near $f_m/f_r = 0.408$, whereas $f_n = 1$ is the series-resonant reference. Light

load produces a pronounced gain peak near the lower boundary; increasing Q suppresses the peak and narrows the high-gain region.

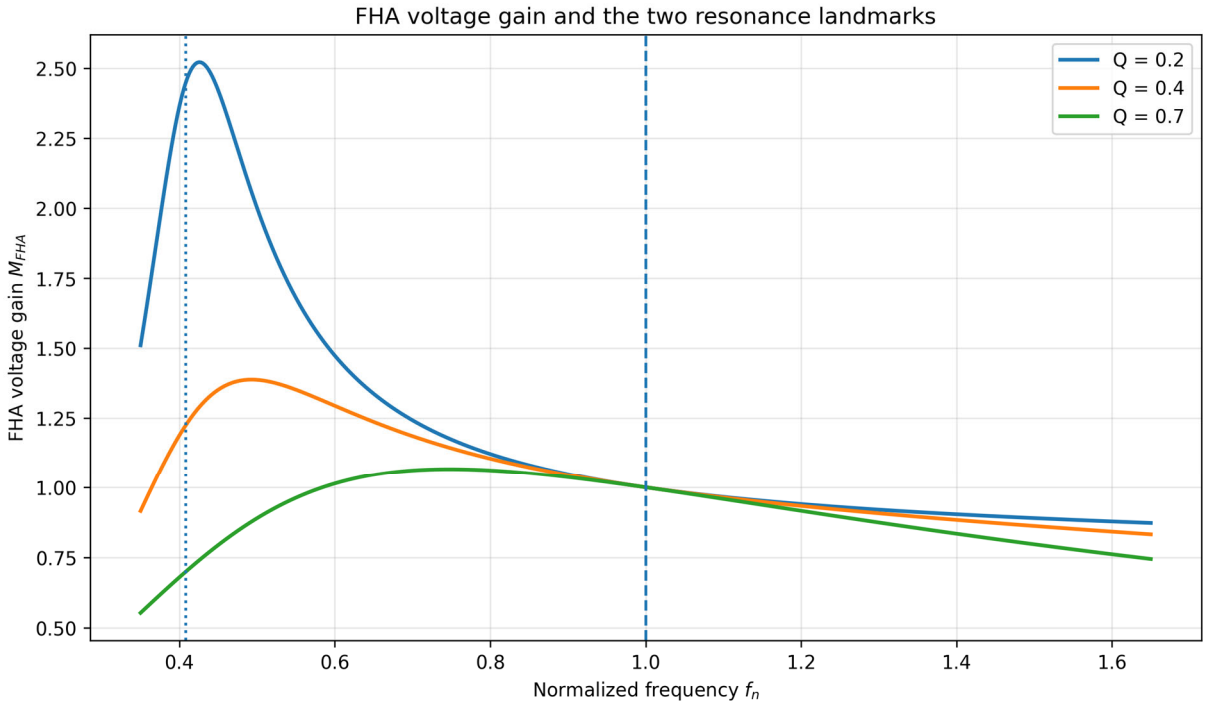

Figure 3. FHA voltage gain for Q = 0.2, 0.4, and 0.7. Vertical lines mark the lower resonance f_m/f_r and the series resonance f_n = 1.

Figure 4 quantifies the reciprocal-frequency gain residual. The residual vanishes at f_n = 1 and decreases as m increases. For Q = 0.4 and the reciprocal pair f_n = 0.80 and 1.25, ε_D is 18.21% for m = 5, 9.02% for m = 10, and 0.089% for m = 1000, which approximates the series-resonant limit. This confirms that finite L_m is the structural symmetry-breaking term in the ideal FHA relation.

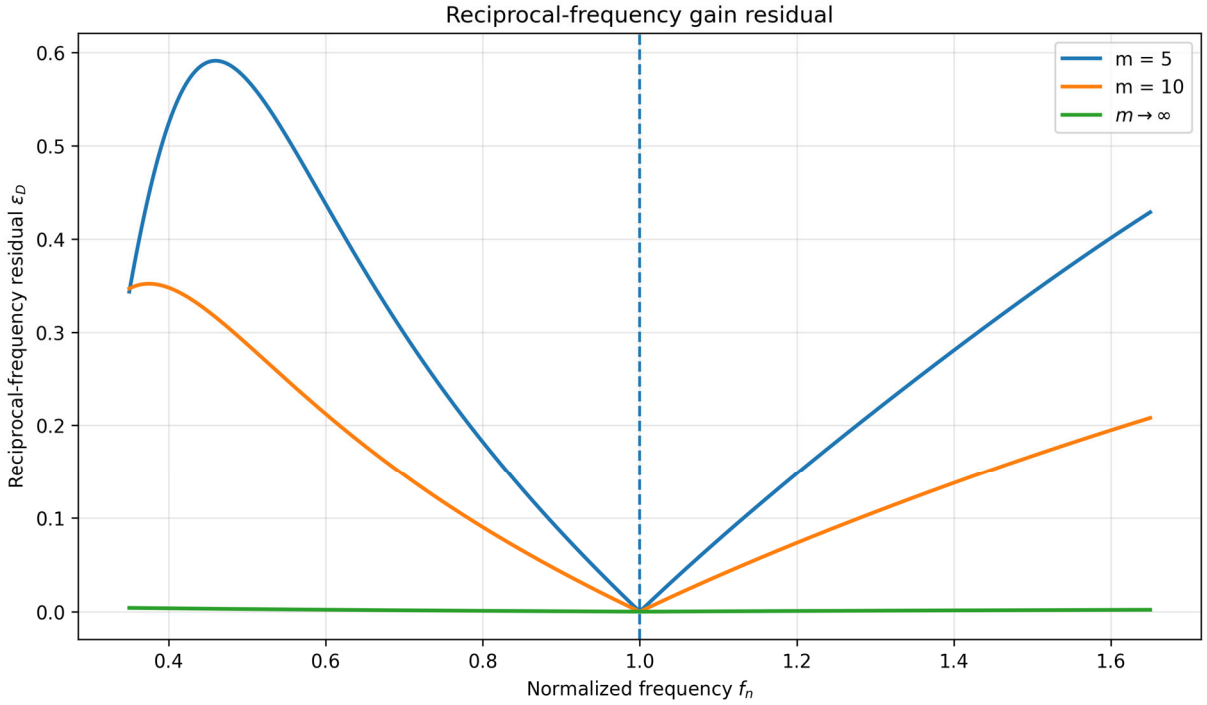

Figure 4. Reciprocal-frequency gain residual for several magnetizing ratios. The vertical dashed line marks the series-resonant point f_n=1, at which the residual vanishes. The large-m series-resonant limit approaches reciprocal symmetry.

The result also clarifies the meaning of extended duality. Sub-resonant and super-resonant points related by $f_n \leftrightarrow 1/f_n$ are not generally gain-identical LLC operating points. They are structured partners whose reactive detuning term is invariant, while the magnetizing branch redistributes gain, current, and commutation behavior.

5.2. Modal Map

Figure 5 presents the categorical canonical modal-family map over the (f_n,Q) plane. Discrete colors are used rather than a continuous scale, and the three selected validation points are marked explicitly. Deep sub-resonant operation is dominated by PON-type words, the intermediate region by PO/O or P words, and super-resonant operation by NP/NOP-related words, under the adopted sign convention.

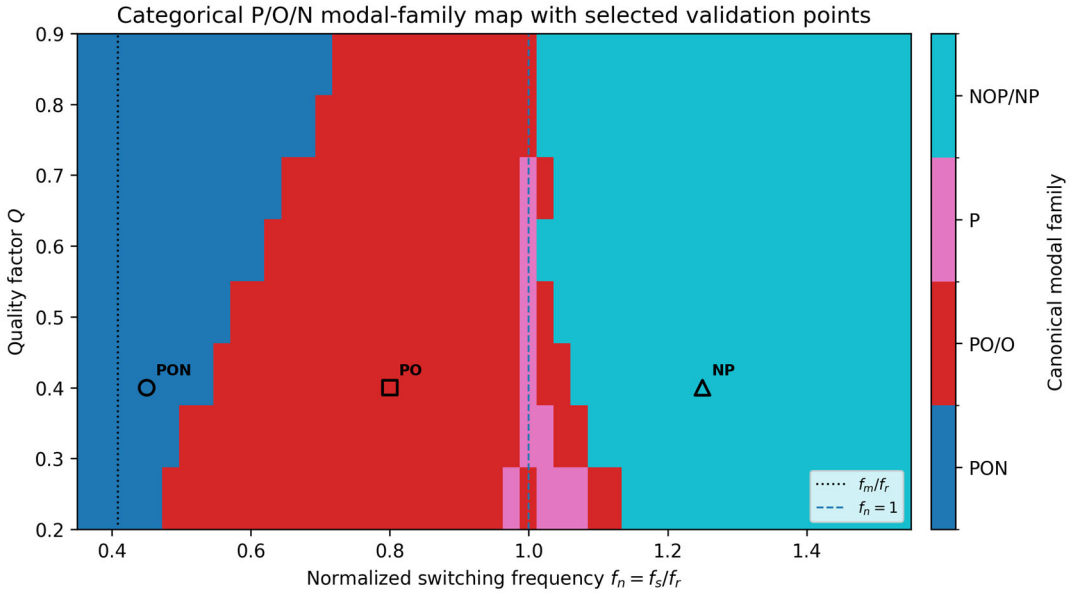

Figure 5. Categorical canonical positive-half-cycle P/O/N modal-family map over the (f_n,Q) plane. The markers identify the PON, PO, and NP points used for detailed low-order and LTspice comparison. The black dotted vertical line marks the lower-resonance reference at f_n=f_m/f_r approx 0.408, and the blue dashed vertical line marks the series-resonant reference at f_n=1. Narrow transition bands depend on event and run-length tolerances and should not be interpreted as universal hard boundaries.

The modal map should not be read as a universal LLC mode atlas. Its boundaries depend on m, load definition, rectifier model, output dynamics, and event tolerances. Its value here is methodological: modal structure is treated as a discrete symmetry coordinate that can be compared alongside continuous gain and stress metrics.

5.3. Representative Mode Waveforms

Figures 6–8 show the three representative low-order operating points. Each plot includes the normalized bridge voltage, resonant and magnetizing currents, resonant-capacitor voltage, and the canonical mode code. The corresponding gain and output-power results are summarized in Table 4, whereas the absolute RMS current and voltage stresses are reported in Table 5.

At f_n = 0.45, the PON word contains an internal secondary-current polarity reversal. It produces the highest gain, current stress, capacitor-voltage stress, and circulation cost of the three points. The oriented ZVS charge margin is negative, predicting failure of ZVS under the adopted dead-time model.

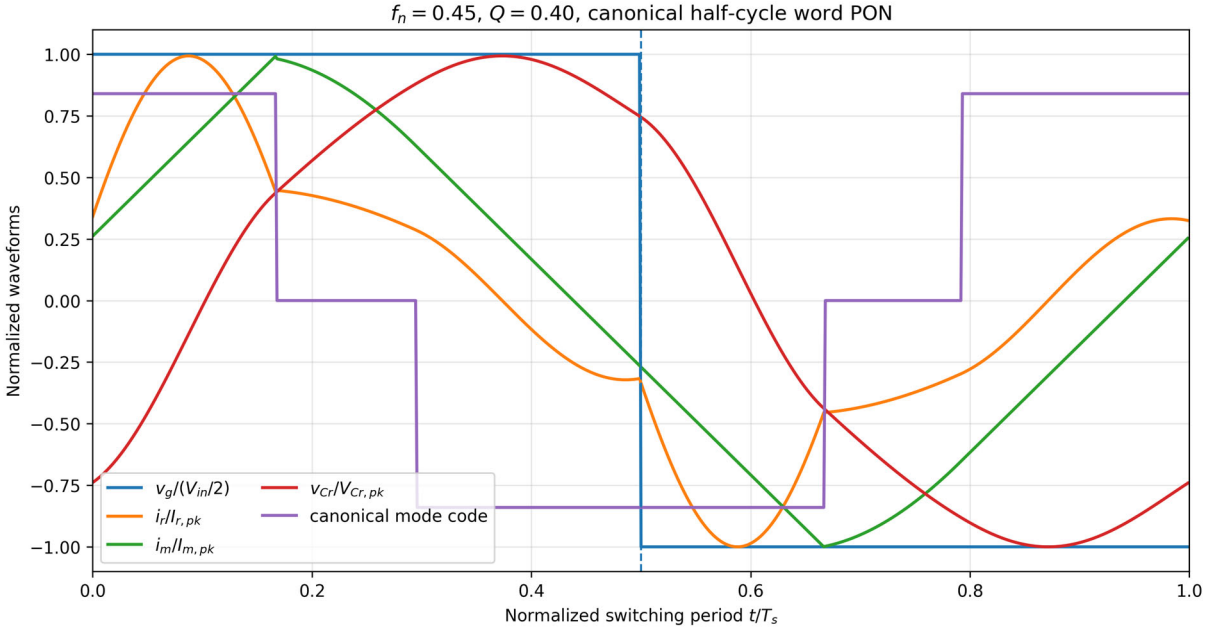

Figure 6. Representative near-lower-resonance PON operation at f_n = 0.45 and Q = 0.40 from the frozen low-order solver.

At f_n = 0.80, the PO sequence is simpler. The two half-cycles have identical canonical modal words, $\varepsilon_H = 1.59 \times 10^{-4}$, and a positive ZVS charge margin. Compared with the PON point, RMS current and circulation are substantially lower.

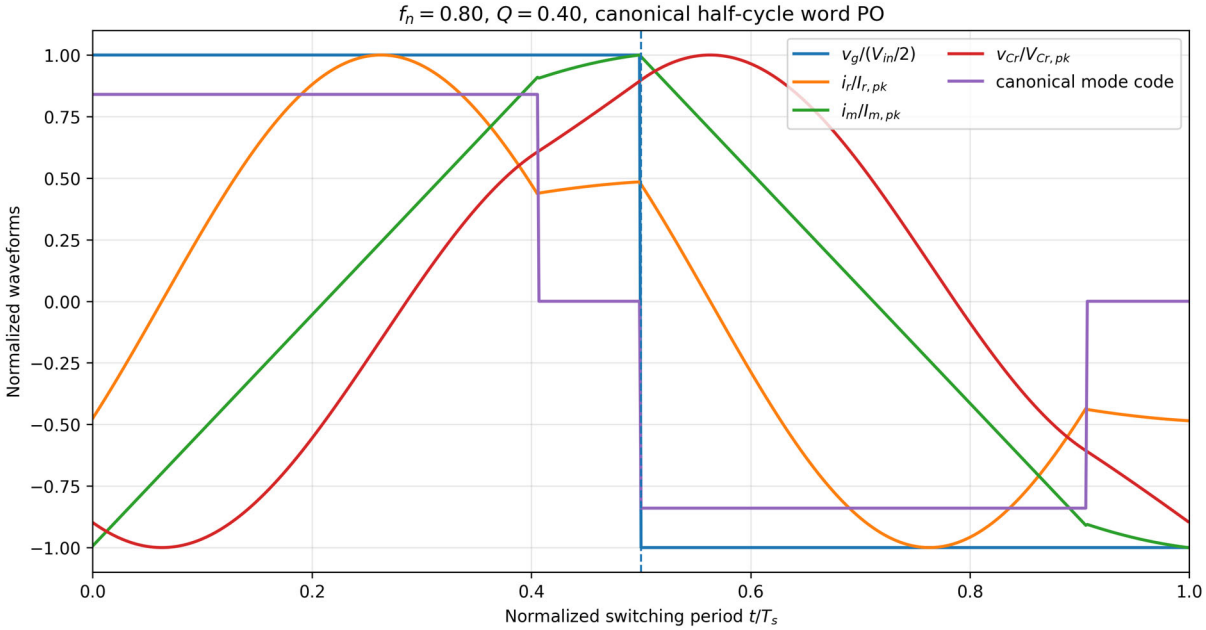

Figure 7. Representative PO operation at f_n = 0.80 and Q = 0.40 from the frozen low-order solver.

At f_n = 1.25, the canonical word is NP. The tank is super-resonant, the gain is below unity, the ZVS margin remains positive, and the circulation index is the lowest of the three

selected points. Finite L_m prevents this point from being a gain-identical reciprocal image of $f_n = 0.80$.

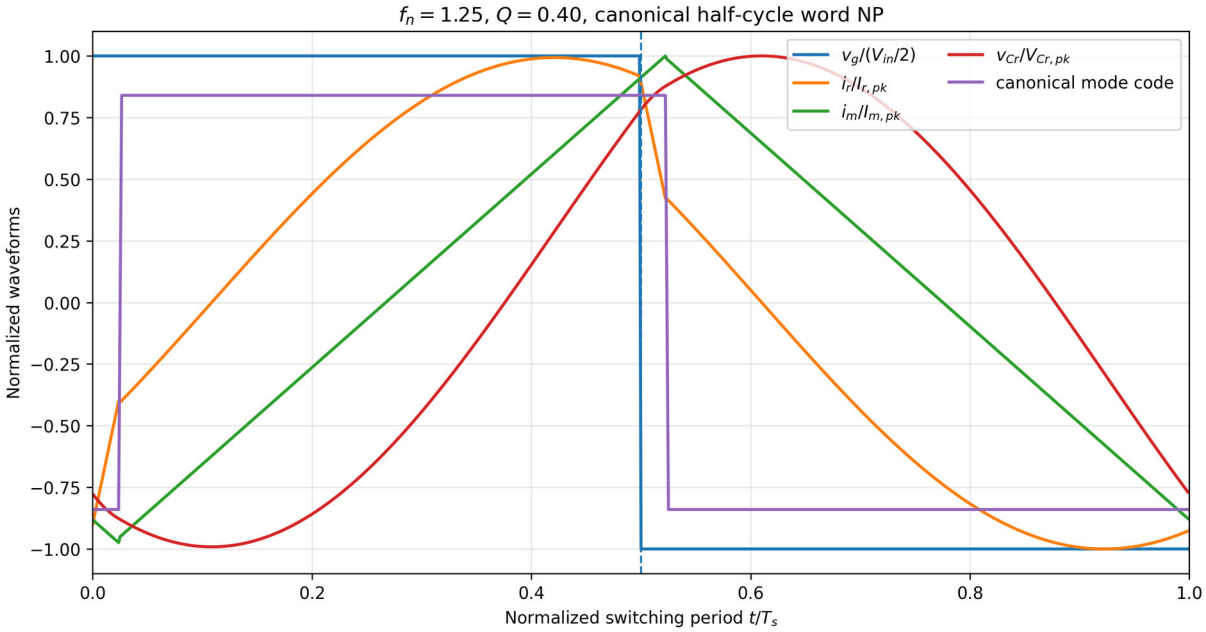

Figure 8. Representative super-resonant NP operation at $f_n = 1.25$ and $Q = 0.40$ from the frozen low-order solver.

Tables 4 and 5 collect the regenerated selected-point results. Relative to FHA, the switched-model gain error is +5.16% at $f_n = 0.45$, +3.62% at $f_n = 0.80$, and −4.18% at $f_n = 1.25$. The absolute $I_{m,\mathrm{RMS}}$ and $V_{Cr,\mathrm{RMS}}$ values make the physical stress differences explicit, while ε_H remains a dimensionless correspondence metric.

Table 4. Regenerated gain and output-power results at the selected operating points.

f_n	Q	Word	M_{TD}	M_{FHA}	Gain Error	P_{out} (W)
0.45	0.40	PON	1.42106	1.35133	+5.16%	694.71
0.80	0.40	PO	1.14420	1.10428	+3.62%	450.38
1.25	0.40	NP	0.88149	0.91996	−4.18%	267.31

Table 5. Absolute stress, event-aligned symmetry, modal, commutation, circulation, and reciprocal-detuning descriptors at the selected operating points.

f_n	$I_{r,\mathrm{RMS}}$ (A)	$I_{m,\mathrm{RMS}}$ (A)	$V_{Cr,\mathrm{RMS}}$ (V)	ε_H	S_{mode}	m_{ZVS}	κ_{circ}	ε_D
0.45	5.546	2.953	448.52	8.33×10^{-5}	1.000	−5.262	2.636	0.5898
0.80	2.790	1.317	129.39	1.59×10^{-4}	1.000	+1.540	0.946	0.1821
1.25	1.967	0.679	58.86	2.06×10^{-4}	1.000	+2.230	0.759	0.1821

5.4. Half-Cycle Symmetry and Symmetry Breaking

Figure 9 shows the event-aligned ε_H across frequency for the ideal and controlled-perturbation cases. The corrected ideal residual remains near the numerical floor over most of the central range; local peaks occur near mode boundaries. Bridge-amplitude, duty, and rectifier-drop mismatch produce persistent physical waveform mismatch even when average gain changes little.

Table 6 reports absolute residuals over $0.65 \le f_n \le 1.35$. The ideal mean is 8.92×10^{-4}, while the four perturbations raise it to 0.0468–0.0794. Because the corrected ideal value is close to the numerical floor, percentage increases would be misleadingly large; the

revision therefore reports absolute residuals and modal scores. Duty imbalance and the combined case also reduce the mean modal score to 0.908.

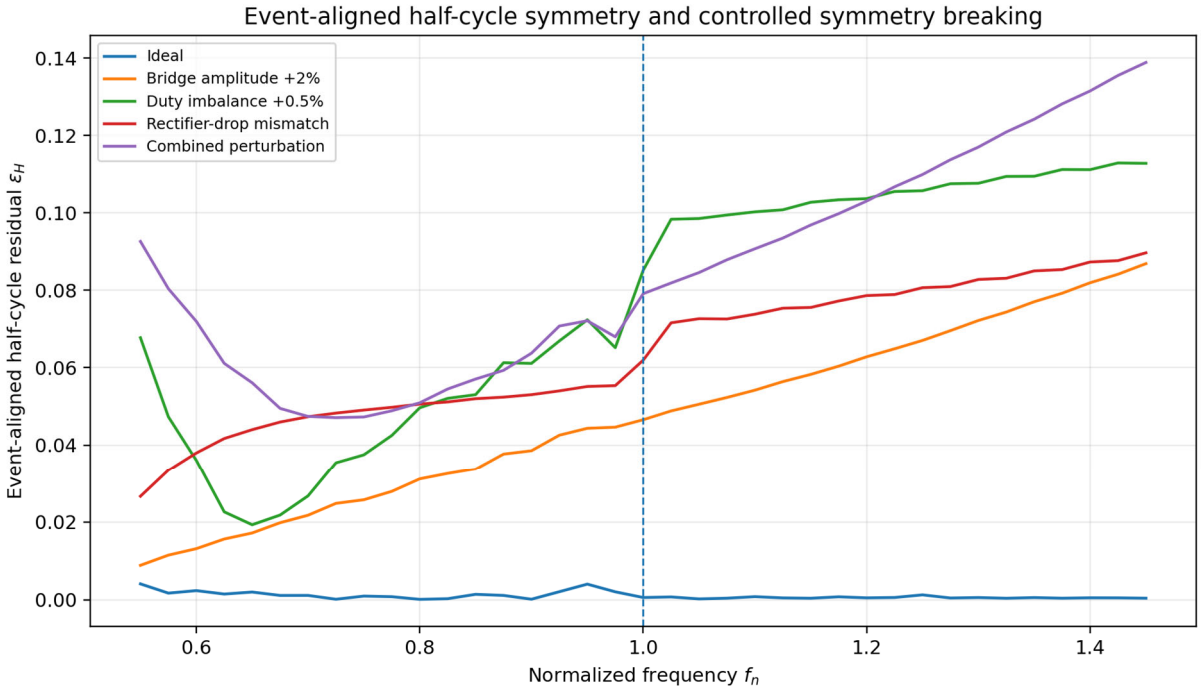

Figure 9. Event-aligned, RMS-normalized half-cycle residual ε_H for the ideal converter and four controlled symmetry-breaking cases at $Q = 0.4$. Actual commanded switching instants are used for alignment. The vertical dashed line indicates the normalized resonant frequency, $f_n = 1$.

Figure 10 summarizes the corrected mean residuals. The metric clearly separates the balanced numerical baseline from implementation asymmetry without relying on unstable percentage ratios to a near-zero denominator.

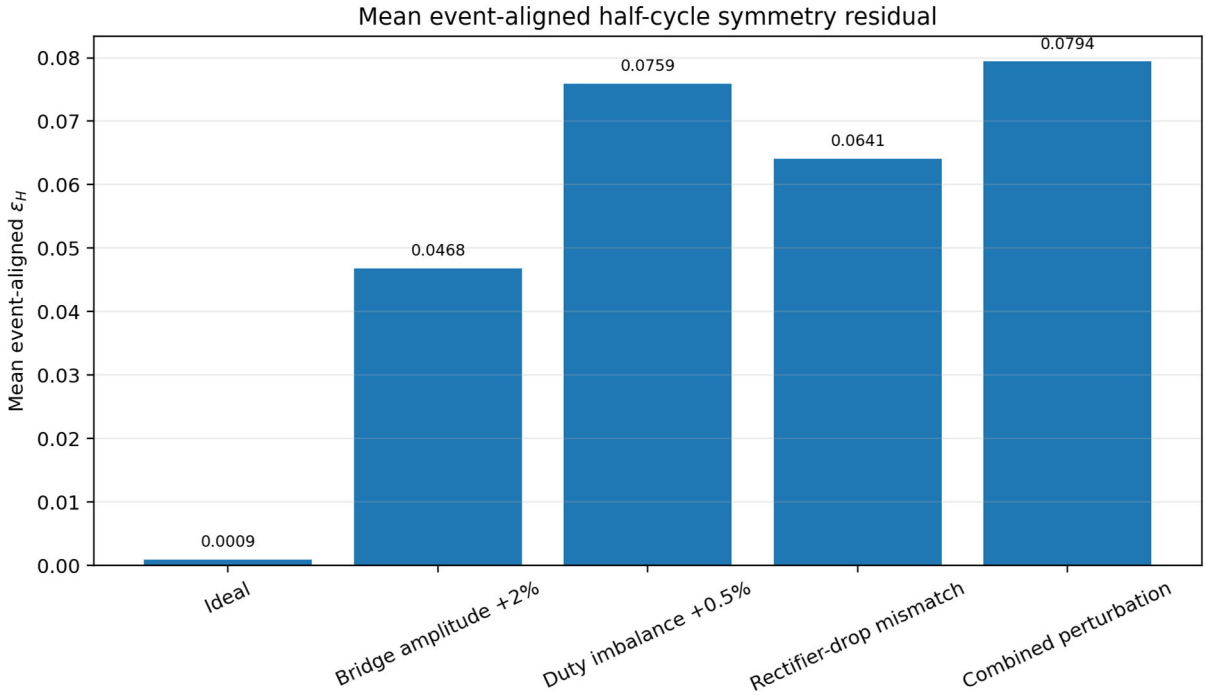

https://doi.org/10.3390/sym18091561

Figure 10. Mean event-aligned half-cycle residual for the ideal and controlled-perturbation cases over $0.65 \le f_n \le 1.35$ at $Q = 0.4$.

Table 6. Event-aligned half-cycle symmetry and modal-score sensitivity to controlled implementation asymmetries over $0.65 \le f_n \le 1.35$ at Q = 0.4.

Scenario	Mean εH	Max εH	Mean Smode	Mean MTD
Ideal	0.000892	0.004013	0.9828	1.03093
Bridge amplitude +2%	0.046803	0.077050	0.9828	1.03094
Duty imbalance +0.5%	0.075948	0.109474	0.9080	1.03094
Rectifier-drop mismatch	0.064102	0.085016	0.9483	1.02598
Combined perturbation	0.079404	0.124219	0.9080	1.02546

Separate Settling, Time-Step, and Event-Location Audits

Figure 11 separates the three numerical effects that were conflated in the original supplementary package. At 720 steps/cycle, M_{TD} converges with cycle count and every selected modal word is retained. At 1300 cycles, the gain approaches a stable value as the fixed-step grid is refined, whereas fixed-grid ε_H is non-monotonic because a discrete sample may move across a rectifier event. The adaptive event-aligned metric removes this aliasing mechanism.

After 20 adaptive polishing cycles, varying the normalized interpolation grid from 600 to 1800 points changes ε_H by only 2.38×10^{-8}, 2.00×10^{-9}, and 4.28×10^{-8} at $f_n = 0.45$, 0.80, and 1.25. The selected 1300-cycle runs satisfy $\varepsilon_{PSS} < 0.01$; extending them to 1800 cycles changes M_{TD} by 0.0036% at $f_n = 0.45$ and negligibly at the other two points. The previous large ideal residuals are therefore identified as fixed-grid/event-alignment artifacts rather than physical symmetry breaking.

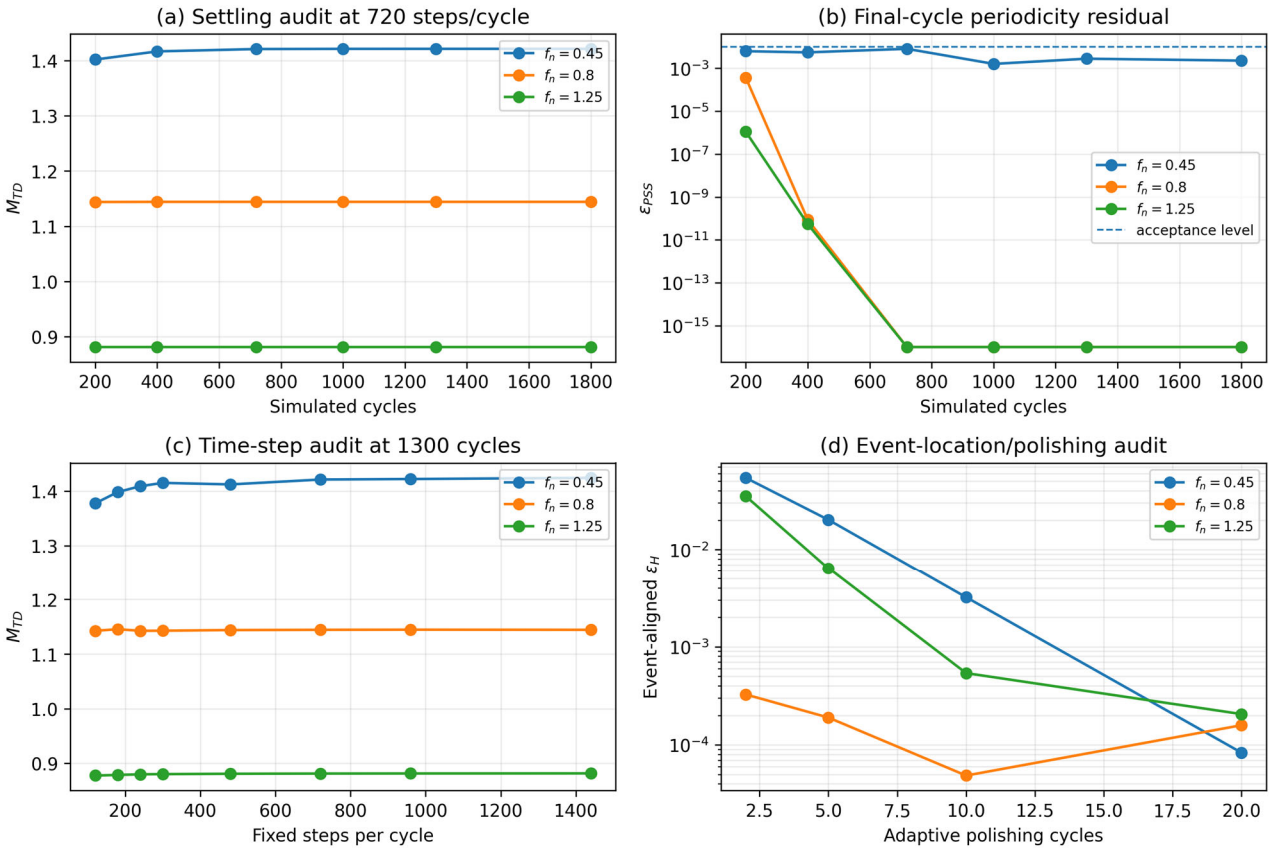

https://doi.org/10.3390/sym18091561

Figure 11. Separate numerical audits: (**a**) gain versus settling cycles at fixed resolution; (**b**) final-two-cycle periodicity residual; (**c**) gain versus fixed steps per cycle at fixed duration; and (**d**) adaptive event-aligned ε_H versus polishing cycles. The event-aligned values are reported as the headline symmetry metric.

5.5. Commutation Margin, Circulation Cost, and Design Corridor

Figure 12 shows the signed screening-level ZVS charge margin. Deep sub-resonant operation fails the adopted charge criterion for all three shown loads. The margin becomes positive closer to resonance, and the super-resonant region provides increasing charge reserve. The exact crossing depends strongly on Q because load changes the resonant-current amplitude.

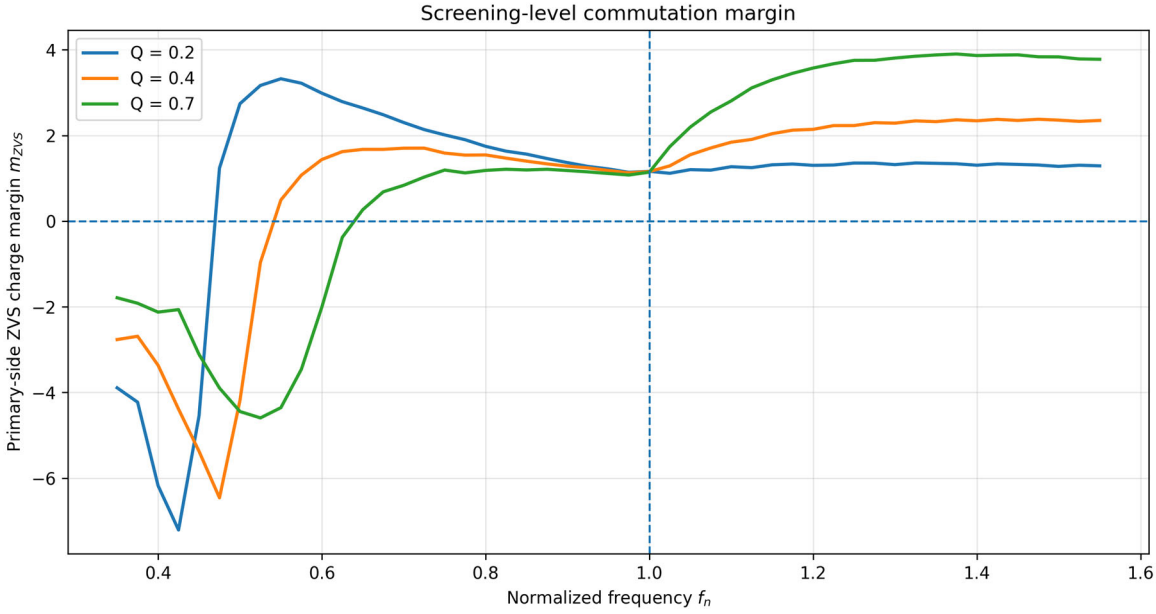

Figure 12. Screening-level transition-oriented ZVS charge margin based on signed resonant current, 200 ns dead time, constant C_{oss} = 200 pF/device, and the required half-bridge commutation charge. The vertical and horizontal dashed lines indicate f_n=1 and the zero-ZVS-margin boundary, respectively.

Figure 13 shows the cost side of the same design decision. At Q = 0.4, κ_{circ} decreases from 2.636 at f_n = 0.45 to 0.946 at f_n = 0.80 and 0.759 at f_n = 1.25. Thus, the high-gain near-lower-resonance regime pays for its gain through increased reactive energy and current stress.

https://doi.org/10.3390/sym18091561

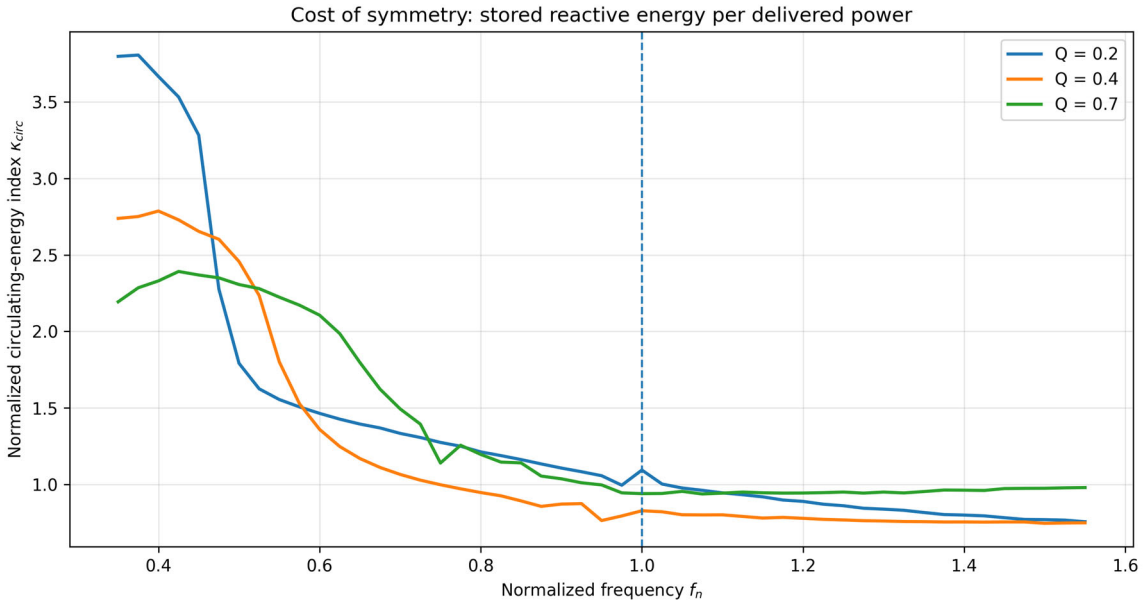

Figure 13. Normalized circulating-energy cost per delivered active power. The metric is intended for relative screening and is not a replacement for a component-resolved loss model. The vertical dashed line indicates the normalized resonant frequency, $f_n = 1$.

Figure 14 shows the nominal illustrative corridor. A point is feasible when $m_{ZVS} \geq 0$ and $0.55 \leq M_{TD} \leq 1.65$. It is preferred when, additionally, $\varepsilon_D \leq 0.25$, $m_{ZVS} \geq 0.20$, and $\kappa_{circ} \leq 1.35$. These replaceable thresholds demonstrate how non-equivalent descriptors can be combined; they do not define a universal optimum.

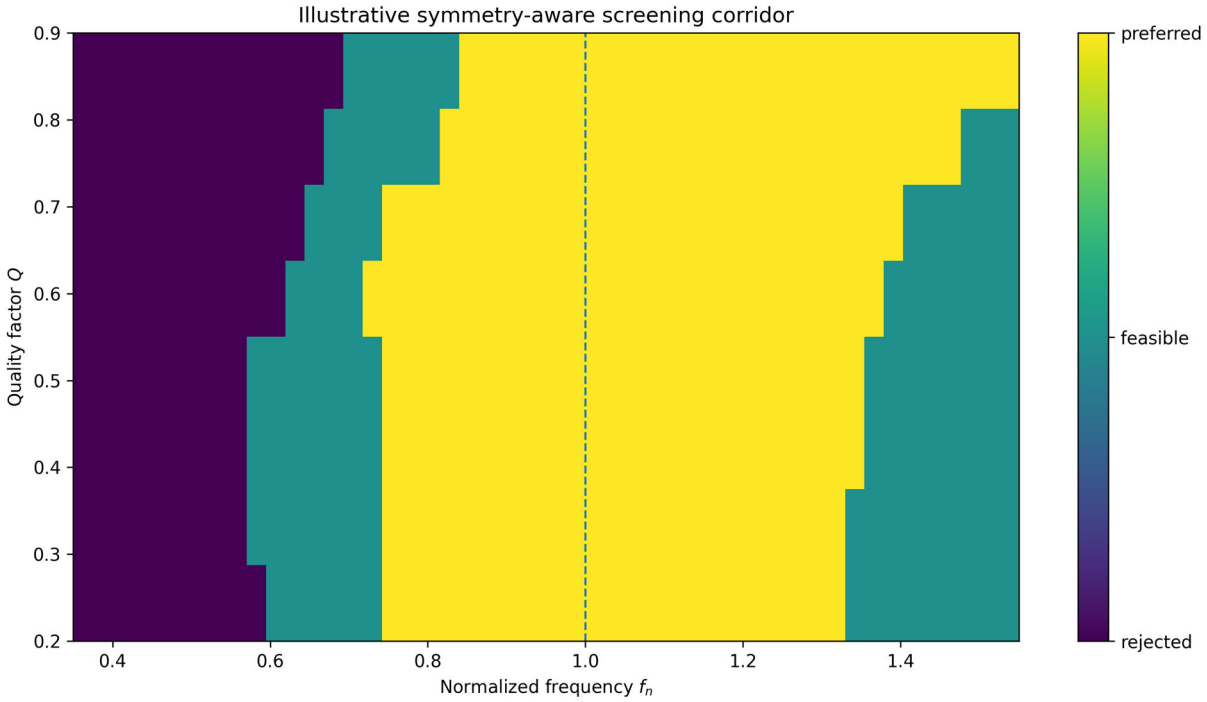

Figure 14. Illustrative nominal screening corridor combining gain feasibility, signed ZVS margin, reciprocal-detuning residual, and circulating-energy cost. The vertical dashed line indicates the normalized resonant frequency, f_n=1.

The preferred fraction is threshold-dependent rather than a model invariant. The nominal criteria retain 53.06% of the grid; the strict and relaxed sets retain 33.67% and

64.29%, respectively (Table 7). The result should therefore be used as a transparent sensitivity example, not as an absolute acceptance rate.

Table 7. Sensitivity of the illustrative screening result to strict, nominal, and relaxed threshold sets.

Threshold Set	Gain Range	ε_D Max	m_{ZVS} Min	κ_{circ} Max	Feasible Fraction	Preferred Fraction
Strict	0.55–1.65	0.20	0.20	1.10	0.7781	0.3367
Nominal	0.55–1.65	0.25	0.20	1.35	0.7781	0.5306
Relaxed	0.50–1.75	0.30	0.00	1.60	0.7832	0.6429

5.6. Tolerance Robustness of the Preferred PO Corridor

The physical-parameter tolerance experiment holds f_s = 80 kHz fixed, so changes in L_r and C_r shift f_r and the actual f_n, L_r and C_r are sampled directly; R_o is sampled as a primitive load quantity; and Q is derived for each draw. Table 8 reports the resulting fifth, 50th, and 95th percentiles.

Table 8. Bounded physical-parameter tolerance results for the nominal 80 kHz (f_n = 0.80) PO corridor (300 deterministic-seed draws).

Metric	5th Percentile	Median	95th Percentile	Interpretation
Time-domain gain M_{TD}	1.1182	1.1439	1.1763	Compact gain spread
ZVS charge margin m_{ZVS}	0.9972	1.5336	2.2143	Positive in all 300 draws
Circulation index κ_{circ}	0.8855	0.9522	1.0248	Moderate cost variation
Reciprocal residual ε_D	0.1552	0.1820	0.2147	Below nominal 0.25 threshold

All 300 draws preserve positive screening-level ZVS margin, canonical agreement between the two half-cycle modal words, and the PO/POP modal family, after run-length compression. This is a local robustness result under independent uniform bounds, not a reliability guarantee. Correlated component variation, temperature coefficients, nonlinear magnetic loss, control quantization, ageing, and nonlinear C_{oss} variation are outside the sampled model. Figure 15 summarizes the resulting distributions of the four principal performance metrics for the 300 deterministic-seed tolerance draws.

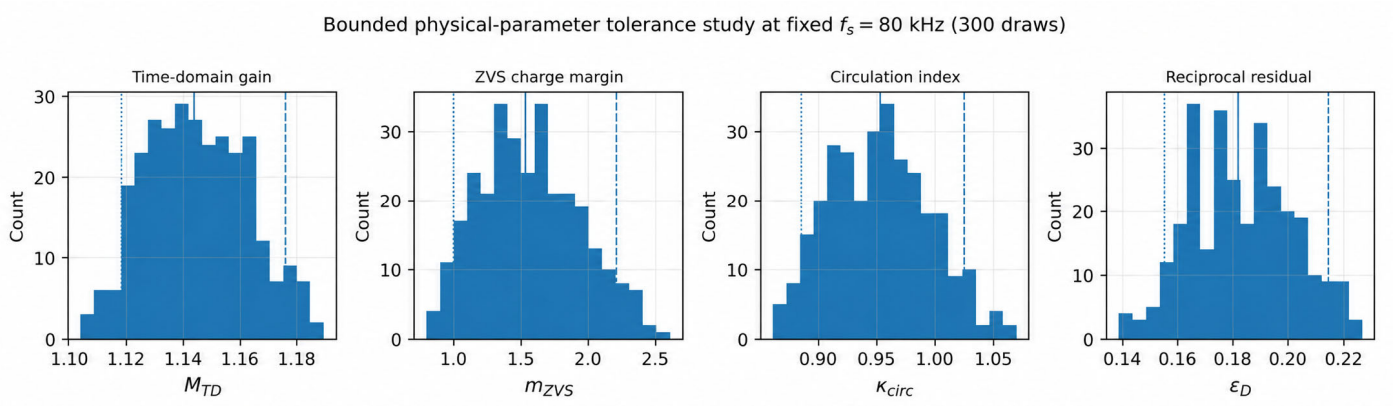

Figure 15. Bounded physical-parameter tolerance study at fixed f_s = 80 kHz: distributions of M_{TD}, m_{ZVS}, κ_{circ}, and ε_D for 300 deterministic-seed draws. Vertical lines mark the 5th percentile, median, and 95th percentile.

The gain and circulation distributions remain compact. The ZVS margin is wider because C_{oss} and dead time enter the charge criterion directly, but its fifth percentile remains positive. The reciprocal residual also remains below the nominal corridor threshold throughout the central 90% interval.

5.7. Independent Parasitic-Aware Circuit-Level Validation

Table 9 and Figure 16 compare the frozen low-order results with the independent LTspice models. The PON, PO, and NP-family (NP or resolved NOP) modal families are reproduced. The gain differences are −0.45%, −3.55%, and −3.11% at f_n = 0.45, 0.80, and 1.25, while the resonant-current difference remains within 4.30%. The larger power difference at the two lower-power points is consistent with the electrical parasitics omitted by the ideal switched model.

At f_n = 0.45, the half-bridge node remains approximately one full bus voltage away from the required rail immediately before each turn-on, confirming hard switching and the negative low-order margin. At f_n = 0.80 and 1.25, residual V_{DS} is below 0.71 V, far below the transparent 20 V practical threshold, confirming completed ZVS in the circuit model. The f_n = 0.80 check uses nonlinear $Q(V)$, while the convergence-safe outer points use charge-equivalent C_{oss}.

Figure 17 shows representative high-side commutations. At f_n = 0.45, the gate command arrives before the node has moved to the upper rail. At f_n = 1.25, the resonant current completes the node transition during dead time, so the commanded turn-on occurs after V_{DS} has collapsed. These waveforms provide the direct circuit-level evidence that the signed screening margin was designed to predict.

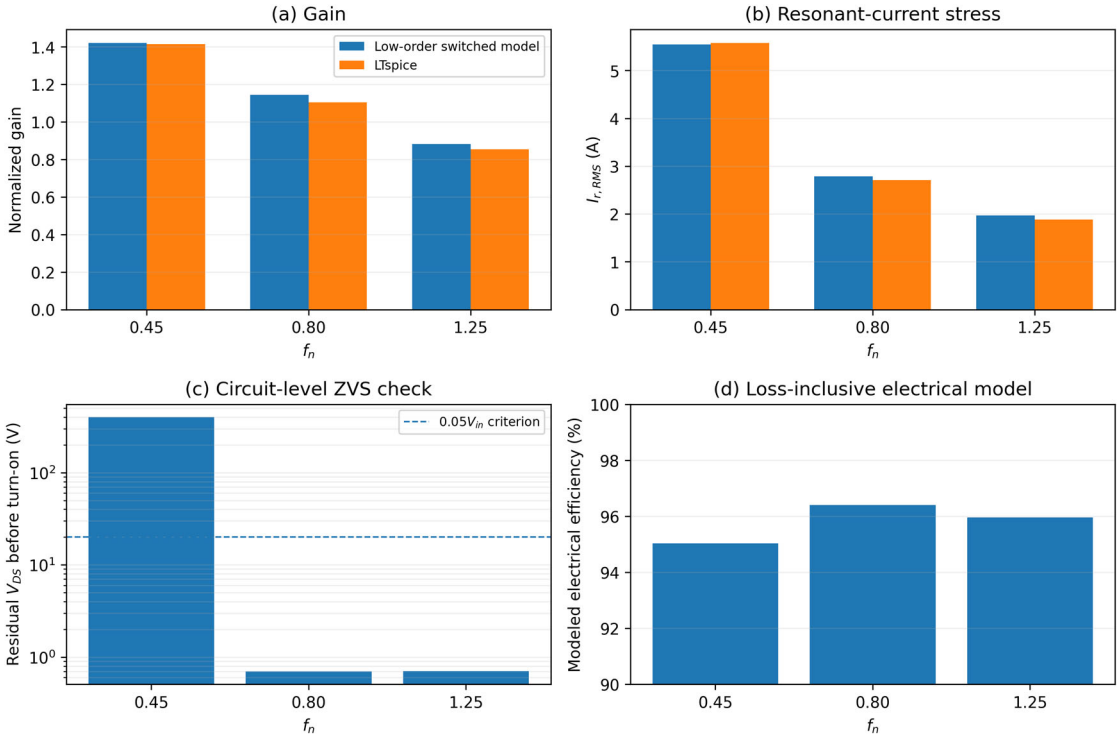

Figure 16. Independent LTspice validation at the three selected operating points: (**a**) normalized gain; (**b**) resonant-current RMS stress; (**c**) residual drain-source voltage immediately before turn-on; and (**d**) modeled electrical efficiency with the explicit electrical losses in Table 2.

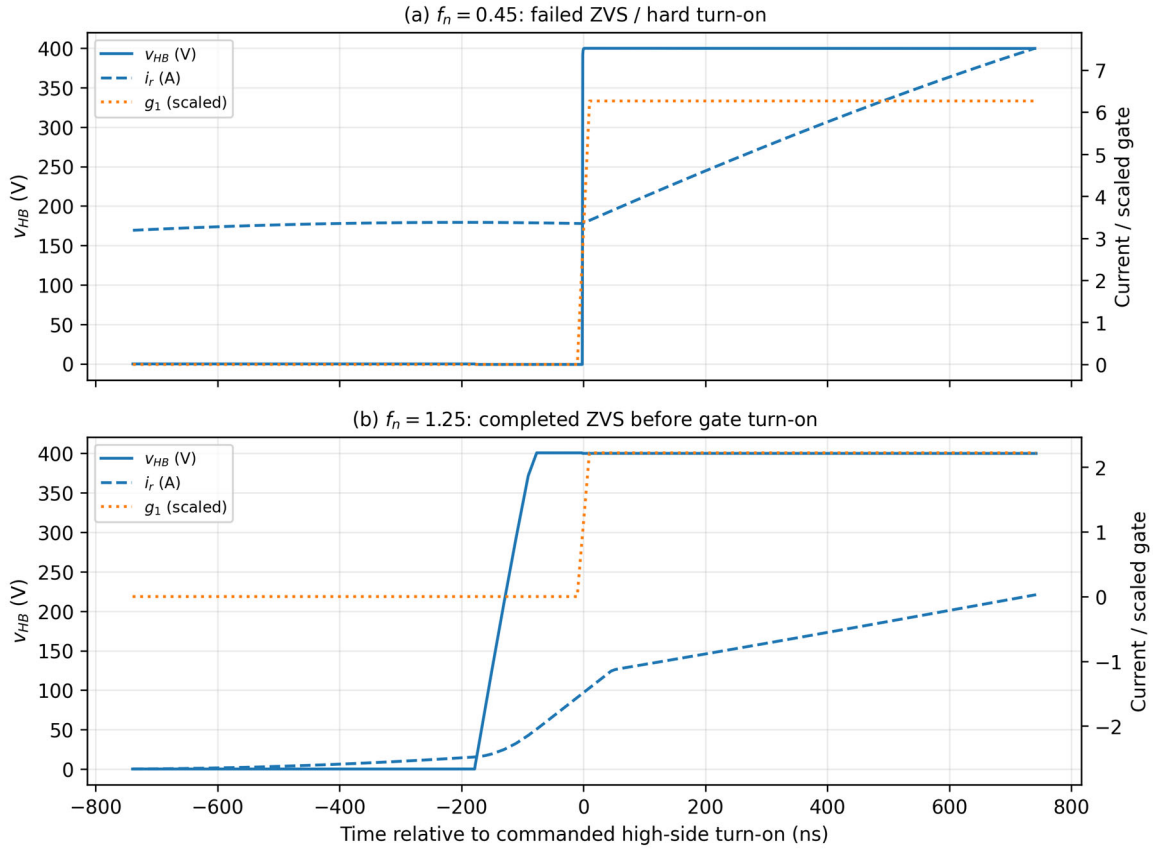

Figure 17. LTspice commutation close-ups referenced to the commanded high-side turn-on: (**a**) f_n = 0.45: failed ZVS; (**b**) f_n = 1.25: the half-bridge node reaches the upper rail before the gate command. Gate amplitude is scaled only for visualization.

Table 9. Independent LTspice spot-check comparison. The η_{elec} is modeled electrical efficiency excluding gate-drive and magnetic-core losses.

f_n	Low-Order/ LTspice Modal Family	M_{TD} Low/LT	Gain Diff.	$I_{r,RMS}$ Low/LT (A)	P_{out} Low/LT (W)	Residual V_{DS} (V)	ZVS	η_{elec} LT
0.45	PON/PON	1.4211/1.4146	−0.45%	5.546/5.577	694.71/688.45	400.78	No	95.04%
0.80	PO/PO	1.1442/1.1036	−3.55%	2.790/2.707	450.38/418.98	0.696	Yes	96.40%
1.25	NP/NP	0.8815/0.8541	−3.11%	1.967/1.882	267.31/250.95	0.707	Yes	95.96%

5.8. Practical Semiconductor Reference and Conditional Loss Budget

Table 10 identifies the practical devices introduced in Table 2 and the parameters actually used. The values have different specified test conditions; they are not treated as one measured operating point. In particular, output-charge and output-energy equivalent capacitances are kept distinct. Replacing either by the small-signal C_{oss} at 400 V would substantially understate commutation charge or hard-turn-on energy.

Table 10. Published device parameters and their use in the separate engineering calculation [34,35].

Element	Part/Source	Published Parameter	Use/Boundary
S_1,S_2	Infineon IPP65R190CFD7A	650 V; typical R_{DS}(on) = 0.159 Ω at 25 °C and 0.356 Ω at 150 °C; V_{GS} = 10 V, ID = 6.4 A.	Channel-conduction reference and temperature sensitivity.
Output charge/energy	Same MOSFET	C_o(tr) = 479 pF; C_o(er) = 46 pF for 0–400 V.	Q_{oss} = 191.6 nC and E_{oss} = 3.68 µJ per device at 400 V.

Gate circuit	Same MOSFET	Q_g = 28 nC; Q_{gd} = 9 nC; plateau 5.7 V; internal gate resistance 10 Ω.	10 V drive and an assumed external 10.2 Ω; charge-based transition-time estimate.
MOSFET body diode	Same MOSFET	V_{SD} = 1.1 V at 6.4 A; Q_{rr} = 0.46 µC at 400 V, 6.4 A, 100 A/µs, 25 °C.	Clamp estimate and an adverse reference-charge scenario, not actual LLC recovery.
D_1–D_4	STMicroelectronics STTH30R02DJF	200 V/30 A; manufacturer loss law: P = 0.77 I_{AV} + 0.006 I_{RMS}^2.	Four individual diodes; two conduct in each secondary conduction interval.
Rectifier recovery	Same diode	Q_{rr} = 140 nC at 125 °C, 30 A, 160 V and −200 A/µs.	Separate reference-charge sensitivity; not a zero-current-switching measurement.

For the piecewise-linear rectifier model, the bridge conduction loss is $2(0.77\, I_{|s|,av} + 0.006\, I_{s,rms}^2)$. The full-cycle MOSFET channel term is $R_{DS(on)}\, I_{r,rms}^2$. A first-order correction removes the channel contribution during the two 200 ns dead-time intervals; a body-diode clamp term is added only for the remaining time after the required charge has been transferred. This correction uses the transition currents, rather than integrating a gate-level waveform.

$$\begin{aligned} P_{\mathrm{D,cond}} &= 2V_{\mathrm{D0}}\, I_{|\mathrm{s}|,\mathrm{av}} + 2r_{\mathrm{D}}\, I_{\mathrm{s,rms}}^2; \\ P_{\mathrm{S,cond}} &\approx R_{\mathrm{DS(on)}}\left[I_{\mathrm{r,rms}}^2 - f_{\mathrm{s}}\, t_{\mathrm{d}}\left(I_{\mathrm{H}}^2 + I_{\mathrm{L}}^2\right)\right] \end{aligned} \tag{25}$$

$P_{D,cond}$ is the total conduction loss of the four-diode bridge and $P_{S,cond}$ is the combined channel loss of S_1 and S_2 (W). The coefficient V_{D0} is the voltage intercept of one diode (V), and r_D is its differential resistance in the adopted linear law (Ω); neither is the complete pair drop $V_{D,s}$. The quantity $I_{|s|,av}$ is the period average of $|i_s|$, not the absolute value of its signed average. The values $I_{s,rms}$ and $I_{r,rms}$ are full-period secondary-current and resonant-current RMS values (A). The factor two in the bridge loss accounts for two series diodes conducting at a time.

$R_{DS(on)}$ is the on-state channel resistance of one MOSFET (Ω). The signed edge currents I_H and I_L are the resonant currents at the upper- and lower-switch commanded turn-ons (A), as defined in Supplementary Material S2. They enter the dead-time correction through their squares. The subscripts H and L identify the high-side and low-side transitions, not half-cycle error or inductance.

A charge-based gate model gives $t_{on} = Q_{gd}(R_{g,int} + R_{g,ext})/(V_{drv} - V_{plateau})$ = 42.28 ns and $t_{off} = Q_{gd}(R_{g,int} + R_{g,ext})/V_{plateau}$ = 31.89 ns. These are first-order Miller-interval estimates at the assumed gate circuit, not transferred datasheet rise/fall times measured at a different drive voltage. A multiplier of 0.5–2 is evaluated in Supplementary Material S2.

In the timing estimates, Q_{gd} is the gate–drain (Miller) charge of one MOSFET (C); $R_{g,int}$ and $R_{g,ext}$ are its internal and assumed external gate resistances (Ω). The drive voltage V_{drv} and Miller plateau voltage $V_{plateau}$ are in volts. The calculated t_{on} and t_{off} are effective overlap intervals in seconds, not the full ON/OFF durations of a switching period.

$$\begin{aligned} P_{\mathrm{off}} &\approx \tfrac{1}{2}V_{\mathrm{in}}(|I_{\mathrm{H}}| + |I_{\mathrm{L}}|)t_{\mathrm{off}}f_{\mathrm{s}}; \\ P_{\mathrm{on,hard}} &\approx [\tfrac{1}{2}V_{\mathrm{in}}(|I_{\mathrm{H}}| + |I_{\mathrm{L}}|)t_{\mathrm{on}} + 2E_{\mathrm{oss}}]f_{\mathrm{s}}; \\ P_{\mathrm{gate}} &= 2Q_{\mathrm{g}}\, V_{\mathrm{drv}} f_{\mathrm{s}} \end{aligned} \tag{26}$$

P_{off} and $P_{on,hard}$ are the combined turn-off and hard-turn-on loss estimates for the two primary switches (W); absolute values of I_H and I_L are used for these overlap estimates. The output-capacitance energy E_{oss} is for one device at V_{in} (J), and Q_g is its total gate charge (C). The quantity P_{gate} is the gate-charge power for both MOSFETs (W). The charge-equivalent $C_{o(tr)}$ and energy-equivalent $C_{o(er)}$ in Table 10 are MOSFET parameters, not the output filter capacitor C_o.

The reference budget assumes negligible reverse recovery. Hard-turn-on overlap and E_{oss} are included when the device-specific signed charge screen fails; they are omitted conditionally when that screen is positive. A positive screen is not proof of completed ZVS for these named devices. A separate adverse scenario applies the 150 °C channel resistance, hard turn-on at both transitions, and the published reference recovery charges. Its additional terms are 2 fs V_{in} $Q_{rr,body}$ and 4 fs $v_{o,av}$ $Q_{rr,rect}$. These scenarios are not confidence intervals or guaranteed upper/lower loss bounds, and the hotter case does not re-solve temperature-dependent waveforms.

$Q_{\mathrm{rr,body}}$ and $Q_{\mathrm{rr,rect}}$ are the per-device reference reverse-recovery charges for the MOSFET body diode and output rectifier, respectively (C); $v_{\mathrm{o,av}}$ is the period-mean output voltage (V). The subscript av denotes an average. The body-diode forward drop V_{SD} in Table 10 enters the clamp estimate described in S2. These parameters retain their declared reference-test conditions and are not measured recovery quantities for this converter.

$$\eta_{\mathrm{budget}} = \frac{P_{\mathrm{out}}}{P_{\mathrm{out}} + P_{\mathrm{S,cond}} + P_{\mathrm{D,cond}} + P_{\mathrm{passive}} + P_{\mathrm{transition}} + P_{\mathrm{gate}} + P_{\mathrm{recovery}}} \quad (27)$$

Here, η_{budget} is the dimensionless conditional electrical-budget efficiency, displayed as a percentage in Table 11. Every power in its denominator is in watts. The term P_{passive} collects the stated winding, resistor and capacitor-ESR losses. The term $P_{\mathrm{transition}}$ combines turn-off, conditionally included hard-turn-on and body-diode-clamp contributions; gate-charge power and the scenario-dependent P_{recovery} are accounted for separately. The labels η_{ref} and η_{adv} distinguish the zero-recovery reference and the adverse scenario, not different definitions of efficiency.

The named-device charge margins are −2.753, +0.061, and +0.296 at f_n = 0.45, 0.80, and 1.25. The small positive value at 0.80 is an important qualification: the favorable margin of the original 200 pF screen does not automatically transfer to a different device. The reference budget assigns 94.74%, 95.48%, and 94.01% efficiency to these three points, but the recovery and temperature scenarios show that practical conclusions depend strongly on commutation conditions. The original 95.04–96.40% generic-model LTspice values in Table 9 must not be relabeled as results for these part numbers.

Table 11. Component-resolved loss estimates for the separate practical reference with C_o present and no added input filter. Powers are in watts; η_{ref} is the conditional zero-recovery reference and η_{adv} is the hot, hard-switching, reference-recovery scenario. Neither is measured efficiency.

f_n	P_{out}	S_1+S_2 Cond.	D_1–D_4 Cond.	Passive	Transition Ref.	Gate	η_{ref}	η_{adv}
0.45	686.67	4.924	18.136	10.242	4.815	0.025	94.74%	91.66%
0.80	415.32	1.126	13.272	3.127	2.078	0.045	95.48%	88.08%
1.25	247.05	0.520	9.885	1.288	3.990	0.070	94.01%	77.58%

The 300-draw robustness result in Section 5.6 applies to its stated screening model and parameter distributions. It is not a robustness qualification of the named-device extension, whose charge margins must be interpreted separately.

The nominal average current per rectifier diode and primary RMS current are below the corresponding headline current ratings under their stated conditions. This does not qualify repetitive current peaks or thermal operation. Overshoot, safe operating area, cooling, repetitive pulse limits, component voltage derating, and magnetic-core loss require device-level and thermal verification. In particular, the high-gain point has substantial resonant-capacitor stress. A component part number is not, by itself, a qualification of the converter.

5.9. Terminal Filtering, Input/Output Waveforms, and FFT

Here "with and without filters" is defined as a terminal-filter ablation. The LLC resonant tank is retained in every case because removing L_r, C_r, or L_m changes the conversion principle. The two independent interventions are an added input LC network and the output smoothing capacitor C_o. This definition distinguishes the resonant power-transfer network from terminal smoothing and makes the compared circuits explicit.

At $f_n = 0.80$, F0 has neither terminal filter; F1 has $C_o = 470$ µF with 25 mΩ ESR only; F2 has the input filter only; and F3 has both. The added input filter has $L_f = 20$ µH, $C_f = 10$ µF, 50 mΩ inductor resistance, and 100 mΩ capacitor ESR. Input voltage, load resistance, tank, transformer ratio, frequency and semiconductor conduction parameters are fixed. Output power is not forced equal: removing C_o changes the loading and conduction intervals, and that change is reported rather than hidden.

With C_o retained, adding the input LC filter changes mean output voltage from 52.739 V to 52.731 V and reduces source-current AC RMS from 1.5609 A to 0.03026 A. The output-capacitor comparison without input LC reduces output-voltage AC RMS from 15.9248 V to 0.14440 V, while mean output voltage and delivered power also change. These are results at one declared operating point, not universal filter attenuation or closed-loop-stability guarantees.

The FFT uses 16 complete switching periods, 2048 uniformly spaced samples per period, a coherent rectangular window, and no duplicate endpoint. The DC component is reported separately in Table 12. The one-sided amplitude is $2|Xk|/N$ for non-DC, non-Nyquist bins; the Nyquist bin is not doubled. Frequency spacing is $f_s/16 \approx 5$ kHz. Figure 18 compares the output-voltage waveforms obtained for the four coupled terminal-filter cases. Figure 19 compares the DC-source current waveforms for the four coupled terminal-filter cases.

Table 12. Coupled terminal-filter comparison at 80.000 kHz. Ripple is RMS AC after removal of the DC component. The source current is measured upstream of the optional input LC network.

Case	Input LC	Output C_o	Mean v_o (V)	v_o AC RMS (V)	Source AC RMS (A)	P_{out} (W)
F0	No	No	45.192	15.9248	1.2451	342.82
F1	No	Yes	52.739	0.1444	1.5609	415.32
F2	Yes	No	45.181	15.9459	0.0241	342.76
F3	Yes	Yes	52.731	0.1446	0.0303	415.18

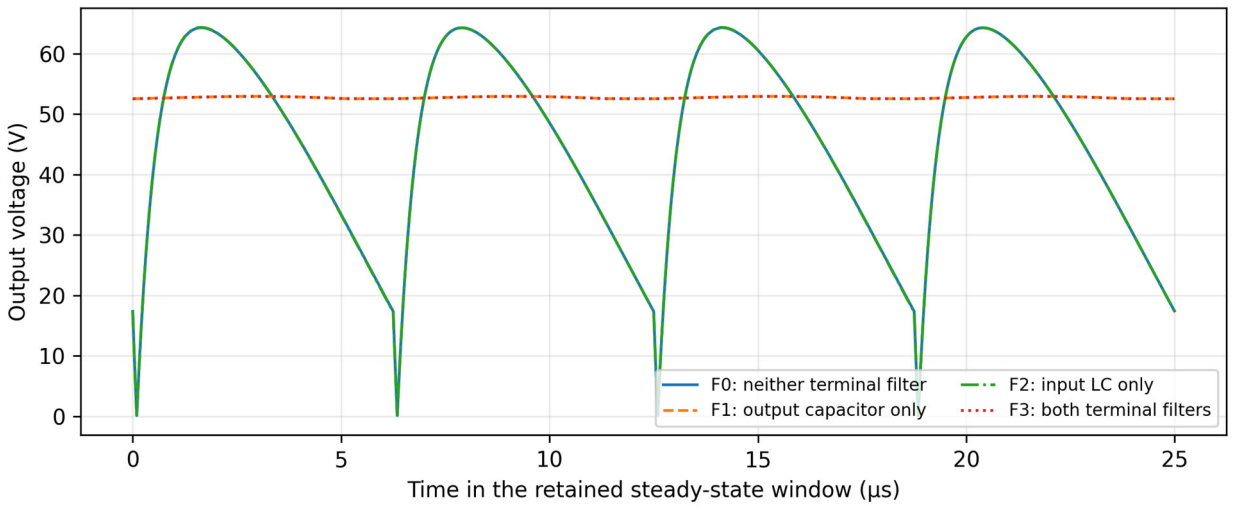

Figure 18. Output voltage in the four coupled terminal-filter cases. Two switching periods from the retained periodic window are shown. The output capacitor changes both ripple and the rectifier loading; mean output values are given in Table 12.

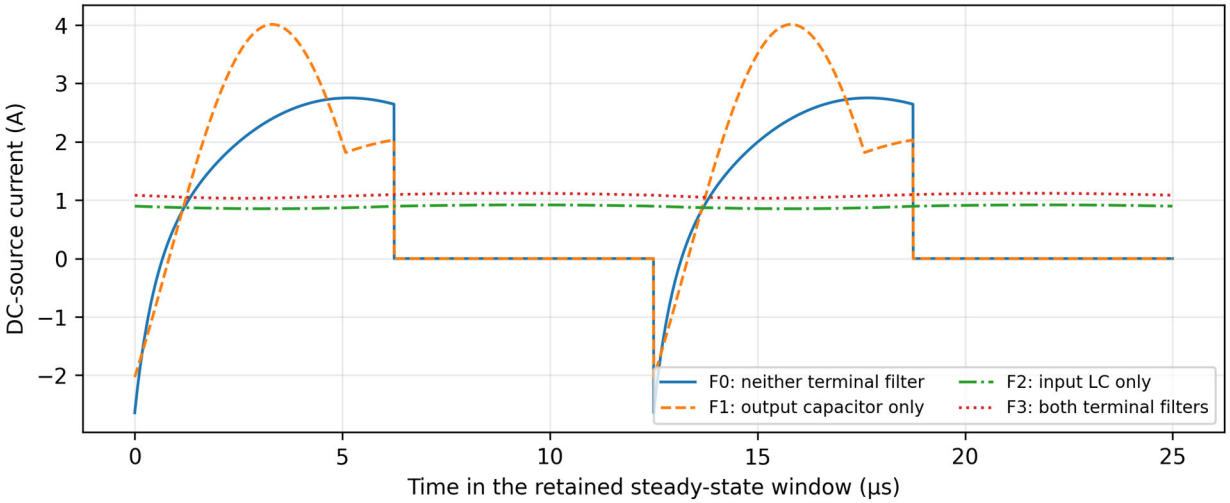

Figure 19. DC-source current for the same four cases, measured upstream of the optional input LC network. F1 and F3 isolate the added input filter while retaining Co. Negative instantaneous values in the unfiltered cases represent energy returned to the ideal DC source.

Figures 20 and 21 display the first 40 switching harmonics in physical peak units. These spectra characterize resolved switching ripple, not conducted-EMI compliance or the high-frequency recovery spectrum omitted by the reduced-order model.

In the spectral expression, X_k is the complex discrete-Fourier coefficient of a mean-removed voltage or source-current record, $|X_k|$ is its magnitude, and N is the number of retained samples, not the negative-secondary mode N. Here, k is a Fourier-bin index, not a bridge transition. Division by N and the stated one-sided factor give physical peak amplitudes in V or A. The source current is measured upstream of the optional filter and is distinct from the resonant current i_r; L_f and C_f denote the added input-filter inductance (H) and capacitance (F).

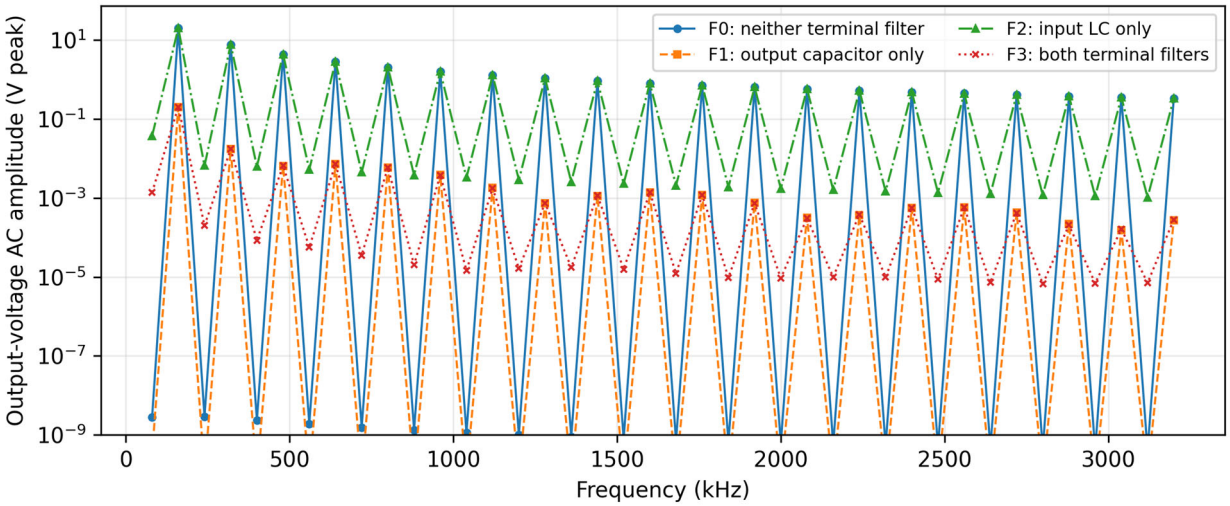

Figure 20. Coherent one-sided output-voltage AC spectra for the four terminal-filter cases. Markers show switching harmonics through 40 fs; amplitudes are peak volts on a logarithmic axis. The DC component is excluded.

https://doi.org/10.3390/sym18091561

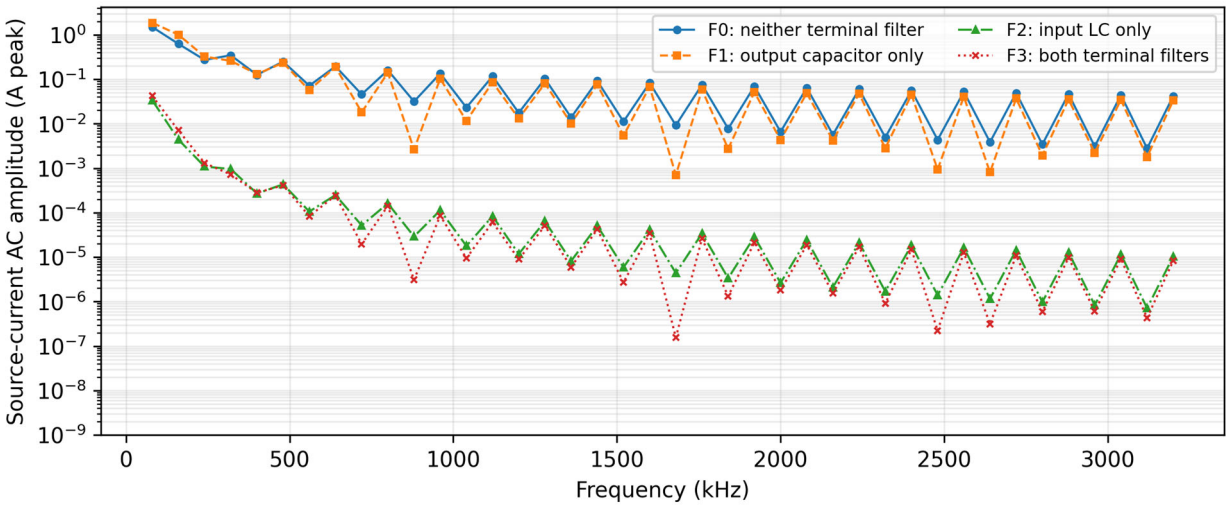

Figure 21. Corresponding DC-source-current AC spectra in peak amperes. Source-side smoothing attenuates the resolved harmonics, but the selected component values do not constitute an optimized or certified EMI filter.

All six additional operating cases satisfy a periodic-state residual below 5×10^{-12} after shooting and verification. The normalized quadrature-based power-balance residual is below 5×10^{-8}. Doubling the sampling density, reducing the maximum solver step and doubling the verification duration results in relative changes in the selected mean/RMS/power metrics of less than 2×10^{-7}. These checks establish numerical consistency at the stated model level, not accuracy of the unmodeled switching or magnetic losses.

6. Discussion

6.1. From Output-Dual Bridge Placement to LLC Modal Symmetry

The family-level classification and the operating-level analysis answer different questions. Classical FHA, time-domain modal analysis, and loss-oriented design already determine gain, stage boundaries, and device stress; the novelty here is not to replace them. It is to place reciprocal detuning, event-aligned half-cycle correspondence, canonical modal pairing, transition-oriented ZVS, and circulating energy in one auditable descriptor set. This exposes trade-offs that are hidden when all favorable behavior is called simply 'soft-switching operation'.

The results illustrate the added design insight. The PON point has exact canonical modal pairing and $\varepsilon_H = 8.33 \times 10^{-5}$, yet it fails both the signed charge screen and the LTspice residual-voltage test and has $\kappa_{circ} = 2.636$. The PO point sacrifices part of the gain but achieves low ε_H, positive ZVS in both models, and much lower circulation. The super-resonant NP point retains ZVS with the lowest current and circulation but lower gain. The framework therefore distinguishes waveform correspondence from commutation feasibility and from efficiency-related stress.

6.2. Finite Magnetizing Inductance as a Structural Duality Breaker

Equation (9) gives a compact interpretation of the LLC branch. The detuning term $Q^2(f_n - f_n^{-1})^2$ is already reciprocal-frequency symmetric. The departure from reciprocal gain symmetry is introduced by the finite magnetizing branch. This creates a direct bridge between the LLC and SRC families: increasing m continuously removes the magnetizing asymmetry and recovers the SRC-like reciprocal response.

This interpretation does not imply that very large m is always desirable. A finite magnetizing current is essential to practical ZVS, and increasing m can reduce the available commutation current at light load. The magnetizing branch is therefore simultaneously a symmetry breaker in the gain domain and a symmetry enabler in the commutation domain. That dual role is one of the most useful design insights obtained from the framework.

This two-level interpretation also clarifies the role of the magnetizing branch. At family level it creates the bridge character of LLC; in the reciprocal-gain relation it is the principal symmetry-breaking term; and at commutation it provides current that can enable ZVS. The same branch can therefore create modal richness, break one symmetry, and enable another. This is more informative than treating L_m only as a gain-shaping element.

6.3. Modal Symmetry Versus Waveform Symmetry

The modal score and continuous half-cycle residual answer different questions. The word is robust to moderate amplitude mismatch while interval order remains unchanged. Event-aligned ε_H detects timing and amplitude differences that preserve that word. Duty imbalance can change both; bridge-amplitude and rectifier-drop mismatch can raise ε_H without changing every modal word. A practical diagnostic should therefore retain at least one discrete and one continuous descriptor.

This separation is especially useful near mode boundaries. A small parameter change can add or remove a short P/O/N interval while integral waveform measures remain continuous. Conversely, an unchanged word can conceal growing current or voltage asymmetry. The corrected event-aligned implementation prevents nominal half-period splitting from misclassifying duty-induced timing difference as a numerical artifact.

6.4. Design Implications

The numerical results make the design principle concrete: the cost of symmetry is redistributed rather than eliminated. Finite L_m breaks reciprocal gain symmetry but supplies commutation current; the high-gain PON regime pays for gain through high stored energy and hard switching; and the PO corridor provides a more balanced combination of gain, ZVS, and stress. The threshold-sensitivity results also show that a preferred corridor must be application-specific rather than universal.

A symmetry-guided LLC design procedure can be summarized as follows:

1. Choose the required gain range and identify the corresponding detuning corridor;
2. Map categorical modal families and avoid narrow transition bands that are sensitive to event or control resolution;
3. Require positive signed ZVS margin under minimum-load and worst-case C_{oss} conditions, then verify retained points by residual-voltage simulation;
4. Screen κ_{circ} together with absolute $I_{r,RMS}$, $I_{m,RMS}$, and $V_{Cr,RMS}$ to quantify reactive stress;
5. Evaluate event-aligned ε_H under bridge-amplitude, timing, and rectifier asymmetry;
6. Perform parasitic-aware circuit simulation and, before deployment claims, hardware validation.

The corridor in Figure 14 is an example of this workflow, not a final design chart. Table 7 demonstrates that its reported area changes materially with the thresholds. Application-specific efficiency constraints and EMI, semiconductor, thermal, magnetic, and control constraints must replace the illustrative values in an actual design.

6.5. Interpretation of the Independent Circuit Checks

The LTspice comparison is intentionally a spot check rather than a second global model. Its agreement in modal family, gain trend, current stress, and ZVS outcome supports the physical interpretation of the low-order descriptors. It also reveals the expected consequence of omitted parasitics: output power is reduced by up to 6.97% and modeled electrical efficiency is 95.04–96.40%. These values should not be compared directly with hardware efficiency because core, gate-drive, temperature, and layout losses are incomplete.

The strongest independent result is qualitative and transition-specific. The negative low-order m_{ZVS} at $f_n = 0.45$ corresponds to approximately 400.8 V residual V_{DS} before turn-on, whereas the positive margins at $f_n = 0.80$ and 1.25 correspond to residual voltage below 0.71 V. Thus, the signed charge equation predicts the direction of the full circuit-level commutation result across all three selected modes.

6.6. Limitations

The practical-device loss budget and terminal-filter comparison extend the engineering interpretation without changing the central symmetry results. They also expose an important limitation: a symmetry-friendly operating point is not automatically robust to a larger device charge requirement, recovery charge, or a different terminal network.

This study remains computational. The main limitations are

- The FHA gain ignores higher harmonics and exact rectifier interval timing;
- The broad switched-model maps use ideal switches and a simplified rectifier clamp;
- The low-order ZVS screen uses constant C_{oss}, although one LTspice point uses nonlinear Q(V) and two use charge-equivalent C_{oss};
- The LTspice netlists include selected electrical parasitics but not calibrated magnetic-core loss, saturation, temperature, or full layout parasitics;
- Dead-time states are represented directly only in LTspice, not in the low-order solver;
- Modal-word extraction retains finite event and run-length tolerances;
- The screening corridor is illustrative and threshold-dependent;
- No laboratory prototype or measured efficiency data are presented.

The named-device extension uses datasheet-informed conduction laws and conditional switching-loss scenarios; it is not a calibrated manufacturer macromodel. Its FFT excludes fast C_{oss} and recovery transients and is not an EMI-compliance test. The no-filter case is explicitly a terminal-smoothing ablation with the resonant tank retained.

Accordingly, the results support mathematical interpretation, numerical reproducibility, and circuit-level trend validation, but not component ratings, reliability guarantees, or measured converter efficiency. A hardware campaign should use vendor devices, measured transformer and capacitor parasitics, calibrated thermal and magnetic loss models, current/voltage probes with stated bandwidth, and repeated uncertainty-aware measurements.

7. Conclusions

This paper developed an LLC-specific symmetry framework that complements established FHA and time-domain modal analysis. At family level, finite L_m gives LLC its multi-resonant bridge character. At operating level, the framework separates reciprocal detuning, event-aligned half-cycle correspondence, canonical P/O/N words, transition-oriented commutation, and circulating-energy cost.

The frozen workflow provides reproducible numerical evidence. The headline ε_H values are regenerated by adaptive event-aligned integration, separate settling/time-step/event audits are reported, the ideal central-range mean residual falls to 8.92×10^{-4},

and the four controlled asymmetries produce absolute mean residuals of 0.0468–0.0794. The 300-run tolerance study derives Q from sampled physical quantities, and the preferred-corridor fraction is reported as threshold sensitive rather than universal.

Independent LTspice checks reproduce the PON, PO, and NP modal-family classifications and the low-order gain within 0.45–3.55%. The NP label is used, consistent with Table 9; a resolved NOP variant remains distinguished, as explained in Section 3.3. The circuit checks confirm hard switching at $f_n = 0.45$ and completed ZVS at $f_n = 0.80$ and 1.25 for the stated models. The practical contribution is a reproducible screening and explanation tool that identifies which symmetry is present, what breaks it, and what current, voltage, and reactive-energy cost accompanies it. Targeted hardware validation and an experimentally informed Pareto-based design methodology remain for a subsequent stage; no prototype performance or measured efficiency is claimed here.

The practical reference illustrates how device charge and semiconductor loss can change the engineering assessment of a symmetry-screened point. The separate terminal-filter study quantifies output-voltage and source-current ripple without removing the resonant conversion network. The resulting loss budgets and spectra remain conditional model-based evidence, while the original LTspice checks retain their separate verification role.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/sym18091561/s1. Supplementary Material S1 contains the frozen symmetry solver, numerical data, convergence studies, generic-model LTspice records, and integrity manifests. Supplementary Material S2 contains the separate device-loss and terminal-filter extension. Supplementary Material S3 contains the retained waveform snapshots and chronology annotations. These materials are provided in the clearly labeled S1, S2, and S3 directories of the accompanying archive. The archived presentation materials relate to Figure 1 in the main text and document the distinction between its fixed-reference state enumeration and the earlier waveform snapshots. Earlier figure layouts are retained for provenance and do not replace the final Figure 1 in the main text. No new numerical or hardware results are introduced by these presentation and packaging updates.

Funding: This work was supported by the European Regional Development Fund under the 'Research Innovation and Digitization for Smart Transformation' Program 2021–2027, Project BG16RFPR002-1.014-0006 'National Center of Excellence Mechatronics and Clean Technologies' The APC was funded by Project BG16RFPR002-1.014-0006.

Data Availability Statement: The computational data, source code, configuration, loss and spectral-analysis exports, and validation reports are included in the combined supplementary archive. The preserved LTspice records belong to the original generic parasitic model; no new manufacturer-macromodel or hardware results are claimed.

Acknowledgments: The research was carried out within Project BG16RFPR002-1.014-0006 'National Center of Excellence Mechatronics and Clean Technologies' The author thanks the reviewers for identifying the numerical-alignment and reproducibility issues addressed in the frozen revision workflow. During preparation of this manuscript, the author used a generative AI assistant for language refinement, software review, figure-layout support, and document formatting. The author independently verified the models, code, calculations, references, figures, and conclusions and assumes full responsibility for the content.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Erickson, R.W.; Maksimović, D. *Fundamentals of Power Electronics*, 2nd ed.; Springer: New York, NY, USA, 2001.

2. Kazimierczuk, M.K.; Czarkowski, D. *Resonant Power Converters*, 2nd ed.; Wiley/IEEE Press: Hoboken, NJ, USA, 2011.
3. Steigerwald, R.L. A comparison of half-bridge resonant converter topologies. *IEEE Trans. Power Electron.* **1988**, *3*, 174–182. https://doi.org/10.1109/63.4347.
4. Yang, B.; Lee, F.C.; Zhang, A.J.; Huang, G. LLC resonant converter for front end DC/DC conversion. In Proceedings of the Seventeenth Annual IEEE Applied Power Electronics Conference and Exposition (APEC 2002), Dallas, TX, USA, 10–14 March 2002; pp. 1108–1112.
5. Fang, X.; Hu, H.; Shen, Z.J.; Batarseh, I. Operation mode analysis and peak gain approximation of the LLC resonant converter. *IEEE Trans. Power Electron.* **2012**, *27*, 1985–1995. https://doi.org/10.1109/TPEL.2011.2168545.
6. Ivensky, G.; Bronshtein, S.; Abramovitz, A. Approximate analysis of resonant LLC DC–DC converter. *IEEE Trans. Power Electron.* **2011**, *26*, 3274–3284. https://doi.org/10.1109/TPEL.2011.2142009.
7. Liu, J.; Zhang, J.; Zheng, T.Q.; Yang, J. A Modified Gain Model and the Corresponding Design Method for an LLC Resonant Converter. *IEEE Trans. Power Electron.* **2017**, *32*, 6716–6727. https://doi.org/10.1109/TPEL.2016.2623418.
8. Wei, Y.; Luo, Q.; Chen, S.; Sun, P.; Altin, N. Comparison among Different Analysis Methodologies for LLC Resonant Converter. *IET Power Electron.* **2019**, *12*, 2236–2244. https://doi.org/10.1049/iet-pel.2019.0027.
9. Deng, J.; Mi, C.C.; Ma, R.; Li, S. Design of LLC Resonant Converters Based on Operation-Mode Analysis for Level Two PHEV Battery Chargers. *IEEE/ASME Trans. Mechatron.* **2015**, *20*, 1595–1606. https://doi.org/10.1109/TMECH.2014.2349791.
10. Hu, Z.; Wang, L.; Wang, H.; Liu, Y.-F.; Sen, P.C. An Accurate Design Algorithm for LLC Resonant Converters-Part I. *IEEE Trans. Power Electron.* **2016**, *31*, 5435–5447. https://doi.org/10.1109/TPEL.2015.2496333.
11. Hu, Z.; Wang, L.; Qiu, Y.; Liu, Y.-F.; Sen, P.C. An Accurate Design Algorithm for LLC Resonant Converters-Part II. *IEEE Trans. Power Electron.* **2016**, *31*, 5448–5460. https://doi.org/10.1109/TPEL.2015.2496179.
12. Geddam, D.; Devaraj, E. Real-time hardware-in-loop implementation of LLC resonant converter at worst operating point based on time-domain analysis. *Energies* **2022**, *15*, 3634. https://doi.org/10.3390/en15103634.
13. De Simone, S.; Adragna, C.; Spini, C.; Gattavari, G. Design-oriented steady-state analysis of LLC resonant converters based on FHA. In Proceedings of the International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM 2006), Taormina, Italy, 23–26 May 2006; pp. 200–207.
14. Duerbaum, T. First harmonic approximation including design constraints. In Proceedings of the Twentieth International Telecommunications Energy Conference (INTELEC 1998), San Francisco, CA, USA, 4–8 October 1998; pp. 321–328.
15. Hong, S.S.; Cho, S.H.; Roh, C.W.; Han, S.K. Precise Analytical Solution for the Peak Gain of LLC Resonant Converters. *J. Power Electron.* **2010**, *10*, 680–685. https://doi.org/10.6113/JPE.2010.10.6.680.
16. Yu, R.; Ho, G.K.Y.; Pong, B.M.H.; Ling, B.W.K.; Lam, J. Computer-Aided Design and Optimization of High-Efficiency LLC Series Resonant Converter. *IEEE Trans. Power Electron.* **2012**, *27*, 3243–3256. https://doi.org/10.1109/TPEL.2011.2179562.
17. Hinov, N. An innovative design approach for resonant DC/AC converters, based on symmetry in their operating modes. *Symmetry* **2023**, *15*, 1864. https://doi.org/10.3390/sym15101864.
18. Hinov, N. A unified approach to the analysis of DC/AC converters, based on the study of electromagnetic processes in a series RLC circuit. *Electronics* **2023**, *12*, 983. https://doi.org/10.3390/electronics12040983.
19. Hinov, N. Symmetry and extended duality in resonant DC–AC inverters: Open-input and closed-input operation below and above resonance. *Symmetry* **2026**, *18*, 599. https://doi.org/10.3390/sym18040599.
20. Vuchev, A.; Grigorova, T. Modeling of an optimal trajectory controlled LLC resonant DC–DC converter. In Proceedings of the 14th National Conference with International Participation (ELECTRONICA), Sofia, Bulgaria, 1–3 June 2023; pp. 1–5. https://doi.org/10.1109/ELECTRONICA58875.2023.11173897.
21. Vuchev, A.; Grigorova, T.; Vuchev, S. Analysis of continuous current mode of an LLC resonant DC–DC converter at ZVS. In Proceedings of the 13th National Conference with International Participation (ELECTRONICA), Sofia, Bulgaria, 19–20 May 2022; pp. 1–4. https://doi.org/10.1109/ELECTRONICA55578.2022.9874422.
22. Grigorova, T.; Vuchev, A. A study of the influence of phase-shift controlled LLC converter parameters on the commutating mechanisms in continuous conduction mode. In Proceedings of the XXXIII International Scientific Conference Electronics (ET), Sozopol, Bulgaria, 17–19 September 2024; pp. 1–6. https://doi.org/10.1109/ET63133.2024.10721536.
23. Beiranvand, R.; Rashidian, B.; Zolghadri, M.R.; Alavi, S.M.H. A Design Procedure for Optimizing the LLC Resonant Converter as a Wide Output Range Voltage Source. *IEEE Trans. Power Electron.* **2012**, *27*, 3749–3763. https://doi.org/10.1109/TPEL.2012.2187801.

24. Adragna, C.; De Simone, S.; Spini, C. A design methodology for LLC resonant converters based on inspection of resonant tank currents. In Proceedings of the Twenty-Third Annual IEEE Applied Power Electronics Conference and Exposition (APEC 2008), Austin, TX, USA, 24–28 February 2008; pp. 1361–1367.
25. Fang, X.; Hu, H.; Chen, F.; Somani, U.; Auadisian, E.; Shen, J.; Batarseh, I. Efficiency-Oriented Optimal Design of the LLC Resonant Converter Based on Peak Gain Placement. *IEEE Trans. Power Electron.* **2013**, *28*, 2285–2296. https://doi.org/10.1109/TPEL.2012.2211895.
26. Luo, J.; Wang, J.; Fang, Z.; Shao, J.; Li, J. Optimal Design of a High Efficiency LLC Resonant Converter with a Narrow Frequency Range for Voltage Regulation. *Energies* **2018**, *11*, 1124. https://doi.org/10.3390/en11051124.
27. Hinov, N.; Gilev, B.; Hranov, T. Model-Based Optimization of an LLC-Resonant DC–DC Converter. *Electronics* **2019**, *8*, 799. https://doi.org/10.3390/electronics8070799.
28. Hinov, N. Quasi-Boundary Method for Design Consideration of Resonant DC–DC Converters. *Energies* **2021**, *14*, 6153. https://doi.org/10.3390/en14196153.
29. Hinov, N. Symmetry and Duality in ZCS and ZVS Quasi-Resonant Buck, Boost, and Buck–Boost DC–DC Converters. *Energies* **2026**, *19*, 883. https://doi.org/10.3390/en19040883.
30. Bhat, A.K.S. A unified approach for the steady-state analysis of resonant converters. *IEEE Trans. Ind. Electron.* **1991**, *38*, 251–259. https://doi.org/10.1109/41.84018.
31. Outeiro, M.T.; Buja, G.; Czarkowski, D. Resonant power converters: An overview with multiple elements in the resonant tank network. *IEEE Ind. Electron. Mag.* **2016**, *10*, 21–45. https://doi.org/10.1109/MIE.2016.2549981.
32. Johnson, S.D.; Erickson, R.W. Steady-state analysis and design of the parallel resonant converter. In Proceedings of the 17th Annual IEEE Power Electronics Specialists Conference, Vancouver, BC, Canada, 23–27 June 1986; pp. 154–165. https://doi.org/10.1109/PESC.1986.7415559.
33. Kazimierczuk, M.K.; Thirunarayan, N.; Wang, S. Analysis of series-parallel resonant converter. *IEEE Trans. Aerosp. Electron. Syst.* **1993**, *29*, 88–99. https://doi.org/10.1109/7.249115.
34. Infineon Technologies. IPP65R190CFD7A: 650 V CoolMOS CFD7A SJ Power Device. Final Data Sheet, Rev. 2.1, 22 November 2021. Available online: https://www.infineon.com/assets/row/public/documents/24/49/infineon-ipp65r190cfd7a-datasheet-en.pdf (accessed on 8 September 2026).
35. STMicroelectronics. STTH30R02DJF: 200 V, 30 A Ultrafast Recovery Diode, High Efficiency. DS8872, Rev. 2, February 2023. Available online: https://www.st.com/resource/en/datasheet/stth30r02djf.pdf (accessed on 8 September 2026).
36. Jiao, J.; Guo, X.; Wang, C.; You, X.; Li, K. An Improved Time-Domain Analytical Method for LLC Resonant Converters and Dead-Time Designs for Zero-Voltage Switching. *IET Power Electron.* **2023**, *16*, 1455–1471. https://doi.org/10.1049/pel2.12482.