

Making True Single-Stage EV Charging Practical: A Dual-Active-Bridge On-Board Charger with Bidirectional GaN and Closed-Form Triple-Phase-Shift Control

Sara Bourouga
EMEA APMS – Marketing & Application
STMicroelectronics
Paris, France
sara.bourouga@st.com

Dario Catalano
Automotive Segment
STMicroelectronics
Agrate Brianza (MB), Italy
dario.catalano@st.com

Julio Brandelero
APMS Power & Discrete
STMicroelectronics
Toulouse, France
julio.brandelero@st.com

Abstract—On-board chargers (OBCs) for electric vehicles are, in most production designs, a cascade of four conversion stages: EMI filter, rectifier, power-factor-correction (PFC) stage and isolated DC-DC stage. Each block costs volume, magnetics and efficiency. This paper reports a system-level reassessment of that chain. A comparative review of the PFC front-ends used in industrial OBCs identifies the single-stage dual-active-bridge (DAB) as the only family able to remove both the intermediate DC-link capacitor and the PFC boost stage while preserving galvanic isolation and bidirectional operation. The realisation described here builds that single stage on a DAB whose primary bridge uses monolithic bidirectional GaN transistors (BDS GaN); the control method developed for it is the subject of a U.S. patent application. Rectification, PFC and isolation are merged into one switching cell operating at a fixed 140 kHz. The control contribution is a triple-phase-shift law computed in closed form, by linear arithmetic, from the instantaneous grid voltage, the battery voltage and the power reference. The secondary phase shift forces volt-second balance so the transformer carries no grid-frequency magnetizing current; a primary-side offset injects a reactive current that sustains zero-voltage switching (ZVS) near the grid zero-crossing and lowers distortion; a differential phase shift sets the transferred power. A 7.4 kW prototype reaches a peak efficiency close to 98 %, a power factor above 99.5 % and a grid-current THD below 1.9 %, with a volumetric density above 5.1 kW/L in 1.45 L. At a matched operating point the simulated grid current improves from 15.7 % THD (series-resonant DAB) to 5 % and becomes visibly sinusoidal.

Keywords—on-board charger, dual active bridge, single-stage AC/DC, bidirectional GaN, triple phase shift, power factor correction, THD, ZVS, V2G, STM32

I. INTRODUCTION

The growth of the battery-electric-vehicle fleet keeps three demands on the on-board charger in tension at once: more power, less volume, and tighter grid compliance to IEC 61000-3-2 [2], IEEE 519 [4] and CISPR 25 [3]. The mainstream industrial answer remains a series chain of converters. A first stage rectifies the line and corrects the power factor, usually with a totem-pole PFC. A second, isolated DC-DC stage — a phase-shifted full bridge or a resonant DAB — regulates the battery current. An input EMI filter and output protection complete the unit; a comparative review of these PFC front-ends is given in [8]. The architecture is mature, but it carries redundant magnetics, an intermediate high-voltage DC bus with

its bulk electrolytic capacitor, and several control loops that interact through a shared ground.

The work reported here started from that observation rather than from a converter-level optimisation. Section III revisits the OBC at system level and asks which of the front-end architectures in current industrial use can actually be collapsed into a single isolated conversion stage. The answer that survives the comparison is the single-stage DAB, and the remainder of the paper is concerned with what it takes to make that answer practical in an automotive product.

Monolithic bidirectional GaN devices (referred to here as BDS GaN) change what is feasible. A switch that blocks and conducts in both directions makes it possible to connect the AC line, through nothing more than an EMI filter, straight to the primary bridge of a DAB, and to let that single bridge perform rectification, PFC and isolated power transfer together. The idea has appeared in the academic literature [1], [5], where 99 % efficiency is reported, but those realisations rely on a variable switching frequency, look-up tables (LUTs) and on-line numerical optimisation, none of which sits comfortably on a cost-constrained automotive microcontroller.

This work makes two contributions. First, a complete automotive-oriented single-stage AC/DC OBC built on BDS GaN, with the packaging, planar-magnetics, sensing and MCU choices needed to actually exceed 5 kW/L at 7 kW. Second — and this is the heart of the paper — a triple-phase-shift control that holds a fixed 140 kHz switching frequency and computes its three phase shifts in closed form, without LUTs, off-line optimisation or on-line interpolation, on a single STM32G4.

The remainder of the paper is organised as follows. Section II sets out the grid and EMC standards that bound a connected charger and the harmonic mechanisms they target. Section III reviews the OBC front-end architectures and the four-to-one stage integration that the bidirectional-GaN primary makes possible. Section IV describes the power topology and the bidirectional-GaN switching cell, including the dual-gate (direct/reverse) structure of the primary devices. Section V is the control method: it introduces the three phase shifts and then derives, in turn, the secondary, primary and differential phase shifts and the closed-form law each one obeys. Section VI is devoted to the modulation stage and the implementation levels of increasing capability that translate the phase shifts into the

actual gate signals. Section VII covers the hardware and the planar-magnetics design, and Section VIII the device packaging, the sensing chain and the digital control. Section IX reports the simulation and experimental results, including the current-quality comparison against a series-resonant DAB. Section X discusses the work against prior single-stage realisations, and Section XI concludes.




Fig. 1. Single-stage SS-DAB converter (full system model): 230 Vac / 50 Hz grid, bidirectional primary bridge (Bdir1–Bdir4), leakage inductance, planar transformer, secondary bridge and 300–400 V battery emulator.

II. GRID STANDARDS AND HARMONIC LIMITS

An on-board charger is, from the grid's point of view, a non-linear load: in a conventional design it draws current in pulses and injects harmonics that range from the low-order line multiples up to the tens-of-megahertz region of switching noise [8]. These emissions are what a set of standards is written to bound. On the low-frequency, power-quality side, IEC 61000-3-2 [2] limits the individual current harmonics a connected device may inject and IEEE 519 [4] caps the total harmonic and demand distortion at the point of common coupling; on the high-frequency, electromagnetic-compatibility side, the automotive product is held to the conducted- and radiated-emission limits of CISPR 25 [3]. Left unaddressed, the harmonics these standards target raise losses in transformers and distribution lines, interfere with communication systems, and can excite resonance and overheating of equipment [8].

The low-order harmonics are the harder ones to remove with a filter, because they sit close to the line frequency, and in totem-pole-type front ends they appear mostly as crossover distortion concentrated in the third, fifth and seventh harmonics near the grid zero-crossing [8]. This is precisely the region the present control acts on. The volt-second balance of Section V-B keeps the transformer free of a line-frequency magnetizing component, the primary phase shift of Section V-C recovers ZVS and feeds the grid current through the zero-crossing, and the freewheeling band of Section VI-D removes the timing uncertainty that would otherwise reintroduce crossover distortion. The combined effect is the measured grid-current THD below 1.9 % and power factor above 99.5 % reported in Section IX, and the simulated reduction from 15.7 % THD for a series-resonant DAB to 5 % at a matched operating point. These figures place the converter inside the low-frequency limits of IEC 61000-3-2 [2] and IEEE 519 [4] by design rather than by added filtering.

Two qualifications are stated deliberately. First, the high-frequency conducted-emission limits of CISPR 25 [3] are met at the system level by an input EMI filter; the prototype density of above 5.1 kW/L is reported without a fully volume-optimised filter, and a formal compliance campaign across the CISPR bands is part of the future work. Second, the harmonic figures

above are measured at the converter terminals and simulated for the comparison case, not certified against the full test procedures of the standards. The point we make is that the control method attacks the low-order harmonics at their source, which is the most filter-resistant part of the spectrum, and so it lightens rather than replaces the compliance effort.

III. ARCHITECTURE REVIEW AND STAGE INTEGRATION

A. Which front-end can be collapsed

The starting point of the project was a comparative review of the PFC architectures in current OBC use, with harmonic content taken as the primary figure of merit [8]. Table I condenses the outcome of that review along the two axes that decide whether an architecture can be integrated further: the harmonic and power-factor behaviour actually reported for it, and the structural element that keeps it from becoming a single isolated stage.

TABLE I. OBC FRONT-END ARCHITECTURES CONSIDERED BEFORE TOPOLOGY SELECTION, AFTER THE COMPARATIVE REVIEW IN [8].

Architecture	Reported harmonic / PF behaviour	Structural limit for a compact OBC
Traditional boost PFC	Crossover distortion at the zero crossing; THD $\approx$ 6 %, PF 0.998 with on-time perturbation	Two bridge diodes always in the conduction path; boost inductor and DC link retained
Bridgeless boost PFC	Smaller crossover zone; near-unity PF with predictive current control	Floating input complicates sensing and EMI filtering; DC link retained
Totem-pole PFC	PF > 0.995 and THD < 5 % in CCM at 6.6 kW	Still a separate isolated DC-DC stage; critical zero-crossing hand-over
H-bridge PFC	Cleaner unipolar-PWM spectrum, first band at twice the switching frequency	Second-harmonic pulsation requires large DC-link capacitance or active filtering
Interleaved totem-pole PFC	THD typically below 3 % at rated load	Higher component count; no effect on the $2 \cdot f_{line}$ ripple
Bridgeless Ćuk-derived PFC	THD 3.49 %, PF 0.9994 at 1 kW	DCM peak currents; practically limited to about 1 kW
Single-stage DAB / LLC AC-DC	PF > 99 %, THD 3–5 % at rated power; 2.9 kW/L demonstrated	Reported realisations use variable frequency (350–650 kHz), complicating EMI filtering and magnetics design

Read as a family, the first six entries share the same limitation. They differ in how well they shape the line current and in how much conduction loss they carry, but every one of them still needs a separate isolated DC-DC stage behind it, and with it an intermediate DC bus and a boost inductor. Improving them further is a matter of degree. Only the seventh entry — the single-stage DAB or LLC AC-DC converter — changes the structure, because the isolated converter itself performs the rectification, the power-factor correction and the isolated transfer [8].

Between the two single-stage candidates the DAB was selected. The resonant LLC realisation reported in the literature buys its soft switching with a variable switching frequency swept over 350–650 kHz, which spreads the conducted-EMI spectrum and forces the magnetics to be designed across a band rather than at a point [8]. The non-resonant DAB, by contrast, can be operated at a fixed frequency provided its three phase shifts are managed correctly, and its modulation space is well documented [7]. That is the trade the present work takes on: accept the control burden of the DAB, and remove it by making the control closed-form.

B. From four stages to one

A conventional OBC contains four functional blocks in series. Integrating the PFC into the DC-DC stage reduces the chain to three blocks; replacing the rectifier diodes with active devices reduces it to two; and once the primary switches of the DAB are themselves bidirectional, the rectifier disappears entirely. The grid then reaches the DAB primary through the EMI filter alone. This is the architecture used here, captured in the full converter model of Fig. 1.

Two structural consequences follow. The high-voltage intermediate DC bus disappears, taking with it the bulk electrolytic capacitor and the reliability liability that comes with it; a small film capacitor across the bridge legs is enough for the commutation transients. The PFC boost inductor disappears as well, its function taken over by the transformer leakage inductance (4.7 μH in the prototype, integrated by design). Together these removals free roughly 30 % of the converter volume relative to a conventional OBC and eliminate one of the dominant copper-loss contributors.

The price is on the control side. The converter now sees a 50 Hz, ± 325 V envelope at its primary terminals while the secondary sees a quasi-DC 300–400 V battery. Keeping the transformer flux balanced under those conditions, within every switching period and across the whole grid period, is exactly what the control of Section V is built to do.

IV. TOPOLOGY AND BIDIRECTIONAL-GAN SWITCHING CELL

The power topology is shown in Fig. 1. The primary bridge is a full bridge of four bidirectional switches, drawn as Bidir1–Bidir4. Each bidirectional switch carries two gate signals, a direct control (S_{xx_d}) and a reverse control (S_{xx_r}): the direct device, when off, blocks a positive drain-to-drain voltage and lets current flow under a negative voltage; the reverse device is its mirror image. This dual-gate structure allows one device to block both polarities of the grid voltage and still provide a freewheeling path. The secondary bridge uses unidirectional devices feeding the battery through a small output filter.

The prototype primary uses 650 V / 15 m Ω bidirectional G-HEMT GaN in a top-side-cooled QDPAK package; the secondary uses 700 V / 23 m Ω unidirectional GaN in TOLT, also top-side-cooled. The QDPAK choice keeps the gate-source loop tight while exposing the package tab to the heatsink, which is what allows fast, low-ringing transitions. A 1:1 planar transformer at 140 kHz, with 4.7 μH of integrated leakage and a volume of about 0.243 L, closes the magnetic loop.

The choice of bidirectional GaN on the primary is what makes single-stage operation possible at all: a standard MOSFET or unidirectional HEMT would need a series diode that would erode the efficiency. As a by-product, since every primary switch already conducts both ways, the same hardware supports vehicle-to-grid (V2G) operation without an added power path.

The monolithic BDS device blocks both a positive and a negative drain-to-drain voltage, and a symmetric current-voltage characteristic is obtained when both gate-source voltages are driven positive. Functionally it replaces the common-source back-to-back pair of two single-gate transistors used in many EV topologies, but in a single die. The penalty is area: for a given on-resistance the monolithic bidirectional switch occupies roughly 1.5 times the silicon-equivalent surface of a unidirectional device of the same rating, a trade-off that is repaid by the removal of an entire conversion stage. The second gate is precisely what provides the freewheeling-diode function the single-stage topology needs during the zero-crossing intervals, and it is the structural reason the zero-crossing freewheeling band of Section VI can define a clean freewheeling state instead of forcing a commutation against an ill-defined polarity.

V. CONTROL METHOD

A. The three phase shifts and why all three are needed

A single-stage AC/DC DAB has three independent control variables that must be coordinated within each switching period: a primary phase shift, which sets the effective duty of the primary bridge; a secondary phase shift, which sets the effective duty of the secondary bridge; and a differential phase shift between the two bridges, which is the dominant power-transfer lever.

A conventional approach sweeps the switching frequency and interpolates the three shifts from a look-up table (LUT) built off-line by numerical optimisation [1]. The drawbacks are well known: a variable frequency smears the conducted-EMI spectrum, the LUT is heavy in ROM and inaccurate between grid points, and on-line interpolation costs cycles an automotive Cortex-M4 can barely spare. The method here uses all three shifts at a fixed frequency, each with a distinct and physically motivated role, computed by closed-form linear arithmetic. The control structure is shown in Fig. 2.




Fig. 2. Closed-form triple-phase-shift control. The secondary phase shift tracks $(|V_{dc}/V_{ref}|) \times 0.5$ (saturated); the primary phase shift adds a power-dependent constant C ; the differential phase shift comes from the power-regulation block.

B. Secondary phase shift — volt-second balance

The aim of the secondary phase shift is to keep the transformer from establishing a magnetizing current at the grid frequency. Volt-second balance over one switching period requires the primary applied volt-seconds to equal the secondary applied volt-seconds, $V_{ac} \cdot dt_{prim} = V_{dc} \cdot dt_{sec}$. Holding the primary duty at its maximum and letting the secondary duty track the grid envelope gives the closed-form law

$$Sec.PS = (|V_{ac}| / V_{dc}) \times 0.5, \text{ saturated to } [0, 0.5]. \quad (1)$$

An absolute-value block keeps the duty positive and a saturation block clips it to $[0, 0.5]$. The result, verified in Section IX, is that the magnetizing current carries no 50 Hz envelope — the single observable signature that the flux is correctly balanced in this topology. The expression assumes a unitary transformer ratio; for a non-unitary ratio V_{dc} is scaled by N_s/N_p . The method is at its best when the transformer is designed so that $V_{ac} < V_{dc} \cdot N_s/N_p$, because the secondary phase shift then never clamps at 0.5.

C. Primary phase shift — reactive current and ZVS recovery

Close to the grid zero-crossing the instantaneous power to transfer is small, and the converter tends to lose ZVS while both terminal voltages are near zero — both effects raise the THD. The primary phase shift is therefore set slightly above the secondary phase shift by adding a constant and saturating the sum,

$$Prim.PS = sat[0,0.5](Sec.PS + C). \quad (2)$$

The added offset deliberately creates a reactive current in the bridge that keeps the primary switches in ZVS even when the grid voltage is low, and at the same time lets the cell transfer the requested power through the low-voltage part of the grid period. Both effects pull the grid current back toward a sinusoid. The constant C is chosen with the operating power: near nominal power it approaches 0.5, while at lighter load it is reduced (toward 0.2). Maximum power transfer is reached when the relevant shift approaches 0.25; pushing beyond that deforms the AC current and is avoided.

D. Differential phase shift — power regulation

With unity power factor the grid current is in phase with the grid voltage, so regulating the peak grid current is equivalent to regulating power. An instantaneous-power reference is formed from the input voltage and the current reference; the measured power is subtracted, and the error feeds a PI controller whose output is the differential phase shift. This shift is saturated to $[-0.25, 0.25]$, the range that spans the full power-transfer capability, and it is the variable through which the user power request is finally satisfied.

We note that the control needs only three measured quantities — the instantaneous grid voltage, the instantaneous battery voltage and the power reference. The three phase-shift references it produces are then handed to the modulation stage, which is the subject of the next section.

VI. MODULATION AND IMPLEMENTATION LEVELS

The modulation block is where the three phase-shift references become the physical gate signals that drive the two bridges. It is also the part of the method that scales from a generic converter to the full bidirectional-GaN single-stage OBC, and for that reason it is presented as a modulator plus four implementation levels of increasing capability. Each level is a strict superset of the one before it, so a designer can take only as much of the method as the chosen hardware actually requires. Table II summarises what each level adds and the hardware it targets; the paragraphs that follow describe them in turn.

TABLE II. THE MODULATOR AND THE FOUR IMPLEMENTATION LEVELS OF THE MODULATION STAGE. EACH LEVEL IS A STRICT SUPERSET OF THE ONE BEFORE.

Lvl	Switch type	What it adds	Targets
A	Modulator	Carrier-and-compare PWM from the three phase shifts; ten HRTIM channels, no LUT	STM32G4 timer
B	Any (incl. unidir.)	Core triple-phase-shift control mapped one-to-one onto the gates	Generic DAB
C	Bidirectional	Grid-polarity steering: one PWM pair drives all four switches of a leg	BDS-GaN primary
D	Bidirectional	Grid-magnitude input and a zero-crossing freewheeling band	Robust polarity hand-over
E	Bidirectional	Phase constant C tuned with power ($PS = 0.0244 \cdot Pin - 0.0671$)	Lowest THD per load

A. The carrier-and-compare modulator

The modulator generates one master triangular carrier between 0 and 1 and four slave triangular carriers at the same fixed switching frequency. Each slave's phase origin is set by comparing the master carrier against one of the phase-shift references produced by the control of Section V; in other words, the phase shifts do not modulate a duty directly, they reposition the slave carriers in time. Each repositioned slave is then compared against a fixed duty value (typically 0.5) to produce the raw PWM signals, which are finally mapped onto the gate drives. Because every operation is a comparison against a counter, the whole modulator maps directly onto the high-resolution timer peripheral of a single STM32G4, which issues the ten PWM channels needed by the two bridges without any look-up table or interpolation. The levels below differ only in how the raw PWM signals are turned into the final gate commands.

B. Generic gate mapping for any switch

The first level is the core method of Section V combined with the carrier-and-compare modulator above, with the PWM signals mapped one-to-one onto the gate drives. In this form the method makes no assumption about the nature of the switches and therefore applies to a converter built with ordinary devices, including purely unidirectional ones. It already delivers the volt-second balance of Section V-B and the power regulation of Section V-D; what it cannot yet do is exploit a primary bridge

that conducts in both directions, which is exactly what the next three levels add. This level is therefore the natural reference against which the benefit of the bidirectional-aware modulation is measured.

C. Polarity-steered gating for bidirectional switches

The second level targets a primary bridge of bidirectional switches. The gate signals are derived from the PWM signals and from the polarity of the grid voltage, and only the primary switches see that polarity; the secondary bridge keeps the modulation of the first level. When V_{ac} is positive (bus+) the reverse switches (S_{xx_r}) are held on while the direct switches commute; when V_{ac} is negative (bus-) the roles swap. Within one leg, PWM_x drives S_{x1_d} and S_{x2_r} while PWM_{xb} drives S_{x2_d} and S_{x1_r} , so a single PWM pair controls all four transistors. The leg-output voltage is then proportional to the PWM duty and to the instantaneous V_{ac} , which is exactly what the control of Section V assumes. Fig. 3 shows the four primary gate signals with the bus+ flag, V_{ac} and the resulting PWM over one grid period; the high-frequency PWM is continuous across the half-period hand-over.




Fig. 3. Polarity-steered gating: the four primary gate signals (S_{x1_r} , S_{x1_d} , S_{x2_r} , S_{x2_d}) with the bus+ polarity flag, the grid voltage and the resulting PWM over one grid period.

D. Zero-crossing freewheeling band

The third level keeps everything from the first two and adds the magnitude of the grid voltage as a third decision input. Inside a small band around the zero-crossing — for example from -15

V to $+15$ V — the converter is commanded into a defined freewheeling state rather than being forced to commute against an ill-defined polarity; the bus+ and bus- states themselves are declared only once the grid voltage passes a threshold of a few volts (e.g. ± 10 V). Two practical benefits follow. First, the grid zero-crossing no longer has to be detected with high precision, because the small band absorbs the timing uncertainty; this removes a fragile element that is a common source of crossover distortion in totem-pole front ends. Second, the freewheeling path created in the band protects the bidirectional switches during the polarity hand-over, since no leg is asked to commute a near-zero, sign-ambiguous voltage. The hand-over itself is realised with turn-on delays and simple combinational (NOT/OR) logic, so it adds no measurable computational load. Fig. 4 zooms in on the zero-crossing: as V_{ac} sweeps from roughly $+13$ V to -13 V the active-switching enable (actSW) drops to zero across the band while the secondary switches S_{y1}/S_{y2} and the secondary PWM continue, which is the freewheeling state described above.




Fig. 4. Zero-crossing freewheeling band: zoom at the grid zero-crossing. The active-switching enable (actSW) drops to zero across the ± 15 V freewheeling band while the secondary switches continue.

E. Power-dependent phase-constant law

The fourth level refines either the second or the third by making the phase constant C of Section V-C a function of the requested power, while the primary phase shift is kept saturated between 0 and 50 %. The motivation is that the offset which recovers

ZVS near the zero-crossing is not equally beneficial at all power levels: too large an offset at light load injects unnecessary reactive current and raises the THD, while too small an offset at high load fails to keep the switches soft. Tuning C against power therefore yields the lowest achievable THD at each operating point. On the prototype the optimum steady-state relation between the phase shift and the input power was found, by measurement, to follow a near-linear law of the form $PS = 0.0244 \cdot Pin - 0.0671$ (with PS the per-unit phase shift and Pin the input power in kW). Fig. 5 plots the measured points against this fitted line over the 3–12 kW range; the close agreement is what allows the constant to be implemented as a single multiply-add rather than as a calibrated table, keeping this fourth level within the same closed-form, table-free philosophy as the rest of the method.




Fig. 5. Power-dependent phase constant: measured optimum phase shift versus input power and the fitted near-linear law $PS = 0.0244 \cdot Pin - 0.0671$ over the 3–12 kW range.

The whole control-plus-modulation chain still needs only three measured quantities — the instantaneous grid voltage, the instantaneous battery voltage and the power reference — and runs entirely as linear arithmetic on the high-resolution timer of a single STM32G4, which generates the ten PWM channels driving the bridges.

VII. HARDWARE INTEGRATION

The prototype is split between a mother board — EMI filter, bus film capacitors, the sensing chain (grid voltage and current, tank current, battery voltage and current, temperatures) and the connectors — and two daughter boards, one per full bridge, each carrying the GaN devices, their isolated single-channel gate drivers, the decoupling capacitors and the eight isolated supplies for the floating gate loops. A small MCU board hosts the STM32G4, the programming interface and the communication port. The assembly, with both flavours of daughter board, is shown in Fig. 6.

The transformer is the densest passive of the design. A 1:1 ratio keeps the secondary phase shift away from its 0.5 saturation at the worst-case V_{ac}/V_{dc} , which matters for the control of Section V-B; the planar core measures $81 \times 68 \times 43$ mm (≈ 0.243 L). The $4.7 \mu\text{H}$ leakage was specified to land on the efficiency plateau of the planar transformer design study at full power. The complete converter fits in roughly 1.45 L, i.e. a volumetric density above 5.1 kW/L at the 7 kW operating point. This density is reported without a fully volume-optimised EMI filter.



Fig. 6. 7.4 kW prototype: motherboard, planar transformer and the BDS-QDPAK / UDS-TOLT GaN daughter boards.

VIII. PACKAGING, SENSING AND DIGITAL CONTROL

A. Device packaging and cooling

The loss budget of a single-stage converter is dominated by the primary switches, so the package was chosen as much for its thermal path and its parasitics as for its on-resistance. The primary BDS GaN is housed in a top-side-cooled QDPAK whose exposed tab faces the heatsink; the same family is offered in bottom-side and dual-side-cooled variants, but the top-side option keeps the gate-source loop short and the power loop tight, which is what limits the switching-node ringing and the associated EMI. Low intrinsic package parasitics are the enabler of the fast, clean transitions on which the ZVS behaviour of Section V-C depends; a noisy commutation would otherwise reintroduce the distortion the control is designed to remove.

B. Power-board partitioning

The converter is partitioned to keep the noisy power loops away from the sensing and the digital control. The mother board carries the bus film capacitors and the complete measurement chain, and exposes the power connectors. Each of the two full bridges sits on its own daughter board, which carries the four bidirectional GaN transistors, eight isolated single-channel gate drivers, six decoupling capacitors and the eight isolated power supplies needed to float the individual gate loops, together with its own mounting pattern. Isolating every gate supply is what allows the four transistors of a bidirectional leg to be driven independently — the degree of freedom the polarity-steered gating of Section VI relies on when it holds one set of devices on while the complementary set commutates.

C. Sensing and microcontroller

The control needs only three measured quantities, but the prototype instruments the full chain — grid voltage and current, transformer tank current, battery voltage and current, and the device temperatures — so that the closed-form laws can be validated against measurement. A dedicated MCU daughter board hosts the STM32G4, the programming interface and a communication port for telemetry and tuning. The high-resolution timer of the STM32G4 generates the ten PWM channels that drive the two bridges directly from the carrier-and-compare modulator of Section VI, with no look-up table and no on-line interpolation in the control loop; the interfaces

between the power and the digital domains are confined to the MCU board, which keeps the analog sensing clean.

D. Transformer design target

The 1:1 planar transformer was sized against a 7 kW design target with a volume budget of about 243 mL. Keeping the secondary phase shift away from its 0.5 saturation at the worst-case V_{ac}/V_{dc} fixed the turns ratio, while the 4.7 μH integrated leakage was placed on the efficiency plateau of the design study at full power so that the leakage inductance doubles as the power-transfer element of the DAB and as the replacement for the boost inductor removed with the PFC stage.

IX. RESULTS


A. Simulation

A full-system PLECS model running the proposed control (Fig. 1) was used to size the magnetics and characterise the grid-current spectrum. At 9.5 kW input power, 230 Vac / 50 Hz grid, 400 Vdc battery, the simulated grid-current THD is 1.23 % and the simulated efficiency is 98.4 %. The simulated transformer magnetizing current shows no detectable 50 Hz component, confirming at model level that the closed-form phase-shift computation balances the flux.

B. Current-quality comparison

Fig. 7 contrasts the simulated grid current of the proposed converter with that of a series-resonant-DAB (SRDAB) reference at the same operating point. The proposed control yields a visibly sinusoidal grid current at 5 % THD, against 15.7 % for the SRDAB reference — the direct consequence of the volt-second balance and the ZVS-recovery offset described in Section V. A 5 % grid-current THD already sits inside the distortion budget set by IEC 61000-3-2 [2] and IEEE 519 [4], whereas the resonant reference does not.

(a) Proposed SS-DAB - grid current, THD = 5 %



(b) Prior art (series-resonant DAB) - grid current, THD = 15.7 %




Fig. 7. Simulated grid current: proposed converter (5 % THD, sinusoidal, top) versus a series-resonant-DAB reference (15.7 % THD, bottom) at the same operating point.

C. Experimental setup

The 7.4 kW prototype was connected to a 230 Vac / 50 Hz grid and a 300–400 Vdc battery emulator, with the heatsink cold plate held at 25 °C. Voltages and currents were captured with a Tektronix oscilloscope; the bidirectional-GaN case temperatures were measured with K-type thermocouples bonded to the package tabs.

D. Efficiency, power factor, THD

Fig. 8 reports the measured efficiency and grid-current THD as a function of power for three battery voltages (300, 350 and 400 V) at a fixed 140 kHz. The peak efficiency is close to 98 %, the power factor exceeds 99.5 % at higher power (above 0.999 near full load), and the grid-current THD falls below 1.9 % over the upper operating range — within the low-frequency harmonic envelope of IEC 61000-3-2 [2] and IEEE 519 [4]. The flat top of the efficiency curves above roughly 3 kW indicates that the BDS GaN devices stay in ZVS for the bulk of the operating window.




Fig. 8. Measured efficiency (left) and grid-current THD (right) versus input power for 300 / 350 / 400 V battery, at a fixed 140 kHz.

E. Thermal behaviour

The case temperatures of the primary BDS GaN devices show only a small rise on the 25 °C cold plate across the operating points, and the optimum ZVS region covers powers above roughly 2.5 kW. The soft-switching window of the QPAK package is therefore broad enough to cover a realistic OBC mission profile without thermal de-rating.

F. Waveforms

Fig. 9 shows the experimental capture at 300 V battery: bridge voltages V_{pri} and V_{sec} , transformer tank current I_{trans} (≈ 7.05 A RMS), grid and battery currents (≈ 4.92 A / 4.72 A RMS) and the grid voltage. The V_{pri} envelope reproduces $|V_{ac}|$ cycle by cycle — the signature of the secondary-phase-shift modulation of Section V-B — and the grid current is sinusoidal and in phase with the grid voltage, without the crossover distortion typical of totem-pole-PFC front ends.




Fig. 9. Experimental waveforms at 300 V battery: primary and secondary bridge voltages, transformer tank current, grid and battery currents, and grid voltage.

X. DISCUSSION

A. Position relative to prior single-stage work

The single-stage AC/DC DAB is not new in itself: realisations reaching 99 % efficiency have been reported [1], [5], and the modulation space of the non-resonant DAB has been mapped in detail [7]. What distinguishes the present work is not the efficiency figure but the way the operating point is reached. The cited works rely on a variable switching frequency together with look-up tables and off-line or on-line optimisation; here the three phase shifts are computed in closed form, by linear arithmetic, at a fixed 140 kHz. The practical consequences are concrete: a fixed frequency keeps the conducted-EMI spectrum concentrated and therefore easier to bring inside the conducted-emission limits of CISPR 25 [3], it lets the magnetics be designed at a single operating frequency rather than across a band, and it removes both the ROM footprint of the table and the interpolation cost from the control loop. The method also reduces the active-switch count relative to a two-stage chain, which lowers both cost and converter volume.

B. Why the control transfers across platforms

Because the core method makes no assumption about the switch technology, the same control migrates from a unidirectional-switch platform to a bidirectional-switch platform by adding only the polarity- and magnitude-aware modulation of the second and third levels. The single technical feature that carries most of the benefit is the primary phase shift of Section V-C: it is what widens the ZVS range, lowers the device stress and recovers the current quality near the zero-crossing, and it is the part of the method that the power-dependent law of the fourth level then tunes against the operating point.

C. Compliance and applicability

The measured power factor above 99.5 % and grid-current THD below 1.9 % place the converter comfortably inside the harmonic-emission envelopes that govern grid-connected chargers, namely IEC 61000-3-2 [2] and IEEE 519 [4], while the automotive context is addressed by the conducted-emission limits of CISPR 25 [3]; a full, volume-optimised EMI filter and a formal compliance campaign remain part of the future work. From a system standpoint the architecture is intended for OBC reference designs in the 7 kW class and, by virtue of the bidirectional primary, for vehicle-to-grid operation on the same hardware.

XI. CONCLUSION

The study began as a system-level comparison of OBC front-ends and ended with a converter that removes two of the four conventional stages. Merging a dual active bridge with bidirectional GaN and a fixed-frequency, closed-form triple-phase-shift control — the subject of a pending U.S. patent application — yields an on-board charger that occupies the volume of a single conventional conversion stage while delivering a peak efficiency close to 98 % and a grid-current THD below 1.9 % across the operating range. The three phase shifts each play a defined role: the secondary shift enforces volt-second balance and removes the grid-frequency magnetizing current, the primary shift injects a reactive current that recovers ZVS near the zero-crossing and lowers distortion, and the

differential shift regulates the transferred power. Removing the high-voltage DC bus and the boost inductor lifts the volumetric density above 5.1 kW/L at 7 kW and eliminates the bulk-capacitor reliability liability, and the bidirectional primary makes the converter V2G-capable by construction.

Future work will integrate the EMI filter within the same envelope, port the control to an automotive functional-safety-class MCU, and explore hosting battery-management functions on the same controller that runs the converter.

REFERENCES

- [1] L. Schrittwieser, M. Leibl and J. W. Kolar, “99 % Efficient Isolated Three-Phase Matrix-Type DAB Buck–Boost PFC Rectifier,” *IEEE Trans. Power Electron.*, vol. 35, no. 1, pp. 138–157, Jan. 2020, doi: 10.1109/TPEL.2019.2914488.
- [2] IEC 61000-3-2:2018, Electromagnetic compatibility (EMC) — Part 3-2: Limits for harmonic current emissions. International Electrotechnical Commission, Geneva, 2018.
- [3] CISPR 25:2021, Vehicles, boats and internal combustion engines — Radio disturbance characteristics — Limits and methods of measurement. IEC, Geneva, 2021.
- [4] IEEE Std 519-2022, IEEE Standard for Harmonic Control in Electric Power Systems. IEEE Power and Energy Society, New York, 2022.
- [5] J. Everts, F. Krismer, J. Van den Keybus, J. Driesen and J. W. Kolar, “Optimal ZVS Modulation of Single-Phase Single-Stage Bidirectional DAB AC–DC Converters,” *IEEE Trans. Power Electron.*, vol. 29, no. 8, pp. 3954–3970, Aug. 2014, doi: 10.1109/TPEL.2013.2292026.
- [6] D. Menzi, J. W. Kolar, H. Sarnago, Ó. Lucía and J. E. Huber, “New 600 V GaN Single-Stage Isolated Bidirectional 400 V Input Three-Phase PFC Rectifier,” in *Proc. IEEE Energy Convers. Congr. Expo. (ECCE)*, Nashville, TN, USA, 2023, pp. 6529–6531.
- [7] N. Hou and Y. W. Li, “Overview and Comparison of Modulation and Control Strategies for a Nonresonant Single-Phase Dual-Active-Bridge DC–DC Converter,” *IEEE Trans. Power Electron.*, vol. 35, no. 3, pp. 3148–3172, Mar. 2020, doi: 10.1109/TPEL.2019.2927930.
- [8] S. Bourouga, D. Catalano and J. Brandelero, “A Comparative Analysis of PFC Architectures in On-Board Chargers Pursuing Zero Harmonics,” in *Proc. PCIM Europe*, Nuremberg, Germany, 2026 (poster).