PROMAGTECH
High-Performance Magnetic Solutions

Title: OBC PFC Inductor Selection & Design Guide: From Single–Phase Totem–Pole to Three–Phase Vienna
Meta Description: How do you size an OBC PFC inductor? This guide breaks down AC loss, FeSi powder–core DC–bias roll–off and saturation limits at 65—140 kHz, contrasts single–phase 11 kW totem–pole and three–phase 22 kW Vienna designs, and gives FeSi / High Flux / Sendust material switch criteria.
Keywords: OBC PFC inductor, flat–wire inductor, FeSi powder core, DC bias, totem–pole PFC, three–phase Vienna, onboard charger, core loss, power density
Canonical: https://www.promagtech.com/blog/obc–pfc–inductor–design

OBC PFC Inductor Selection & Design Guide: From Single-Phase Totem-Pole to Three-Phase Vienna

Power inductor engineering for 800 V-platform onboard chargers · ProMagTech Engineering Team

Key Takeaways First

In OBC PFC inductor design at 65–140 kHz, what decides success is never DC resistance (DCR) — it is high-frequency AC loss and the permeability roll-off under DC bias.

If you take only one thing from this article, take the table below. It puts the most common engineering traps right up front:

Key question	One-line answer
What to size on first?	Inductance retention at peak current (L@Ipk) and temperature rise — not zero-bias inductance or DCR.
Where does flat wire win?	At high current density and high frequency, by cutting AC resistance (skin + proximity), not by piling copper to lower DCR.
FeSi powder-core limit?	μ60 FeSi is most cost-effective for single-phase 11 kW; for three-phase 22 kW with heavy DC bias, evaluate switching to High Flux.
Single vs three-phase?	Three-phase carries larger per-leg current and heavier DC bias; saturation margin and thermal rise dominate — do not copy the single-phase design.
The most hidden trap?	Passes on the bench, overheats in the unit — usually from ignoring bias roll-off and the potted-enclosure thermal path.

Note: *All parameters use public industry reference values (standard Bs, Curie temperature, typical μ-Hdc curves, etc.), labeled as reference values to be confirmed per project.*

OBC Topology & the Inductor's Role Boundary

First, clear up a common misconception: the OBC (onboard charger) is the AC slow-charge path, typically 6.6 / 11 / 22 kW. True high–current DC fast charging is handled by the charger–side DC path, bypassing the OBC. So the PFC inductor's challenge is not “how large is the current,” but “how do you control loss and temperature rise in the smallest volume, at a given switching frequency and bias.”

Two mainstream topologies

First, a one–line plain–language note on these two topology names for the discussion that follows — **Totem-pole PFC:** replaces the diodes in the traditional PFC input rectifier bridge with a pair of series high–frequency switches forming a “totem–pole,” which both rectifies and boosts, giving higher efficiency and lower bridge loss — the mainstream for high–efficiency single–phase OBCs; **Vienna rectifier:** a three–phase, three–level topology where each phase clamps the output to a midpoint, yielding three–level operation, low current ripple and low switch stress — a common choice for three–phase medium/high–power PFC. In both, the PFC inductor works as a high–frequency boost inductor, but with different phase count and bias conditions.

Topology	Single-phase totem-pole	Three-phase Vienna / totem-pole
Typical power	6.6–11 kW	11–22 kW
Inductor role	Boost storage / filter inductor	Per-phase boost inductor (×3)
Per-leg current	Higher (one phase carries full power)	Lower per leg, but with heavy DC component
Ripple character	Switching-frequency ripple dominant	Interleaving partially cancels ripple
Main constraint	HF AC loss + thermal rise	μ roll-off under DC bias + saturation margin

Interleaving is the key technique three–phase architectures use to cut input ripple: multi–phase currents are time–staggered so the composite ripple cancels, which **relaxes the inductance requirement per unit** — but each inductor must still independently withstand the DC bias of its own leg current. That is exactly why FeSi powder cores reach their limit in three–phase scenarios.

Loss Decomposition: DCR Is Not the Lead

PFC inductor total loss = winding loss + core loss. Winding loss further splits into DC resistance loss ($I^2 \cdot DCR$) and high–frequency AC loss (skin + proximity). Engineers' intuition often stops at DCR, but at 65—140 kHz the added AC loss is frequently the real driver of temperature rise.

3.1 Why flat wire wins at high frequency and high current

Skin effect pushes HF current toward the conductor surface; proximity effect redistributes current between adjacent conductors and layers. Together they raise the effective AC resistance Rac. Flat wire, with its "wide and thin" cross–section, compresses conductor thickness in the direction perpendicular to the field, **markedly suppressing proximity loss**, while edge–wound construction raises window fill factor and shortens mean turn length. The result: at the same current, flat wire shows lower AC loss and lower temperature rise than equal–area round wire — most of all in PFC inductors with high current density and many layers.

Dimension	Conventional round wire	Flat wire (edge-wound)
Rac/Rdc ratio	Higher (proximity stacks over layers)	Lower (thin section suppresses proximity)
Window fill factor	Limited by round-wire gaps	High (tight fit)
Heat path	Little inter-wire contact	Wide-face contact, better conduction
Rise at high current density	Higher	More controllable
Production consistency	Sensitive to winding tension	Edge-wound tooling is repeatable

One layer deeper: proximity loss depends strongly on the number of winding layers. Classic Dowell analysis shows that in a multilayer winding, inner conductors are penetrated by both their own field and that of adjacent layers, so each added layer raises Rac/Rdc not linearly but with acceleration — the more layers, the worse it gets. For PFC inductors with many turns this means: **simply adding turns to raise inductance comes at the cost of a nonlinear blow-up in AC loss.** Flat wire's value is precisely that its thin section keeps each layer's effective thickness within the skin depth, attacking inter–layer proximity at the source.

A practical criterion: when conductor thickness approaches or exceeds the skin depth δ at that frequency (δ shrinks as frequency rises), AC loss begins to run away. Match the flat–wire thickness dimension to δ in the target band — **sizing the conductor by I^2R alone is the most common starting error in HF inductor design.**

3.2 Core loss: nonlinear in frequency and flux swing

Core loss approximately follows the Steinmetz relation Pv ≈ k·f^α·B^β, where B is the flux swing. **Note: a PFC inductor operates with large DC bias,** so what drives core loss is the HF ripple swing ΔB riding on the line–frequency envelope, not the peak flux. Lowering ripple (higher inductance or interleaving) directly lowers core loss — but higher inductance means more turns and higher winding loss.

Quantify the chain: ΔB is set by the volt–second integral across the inductor, ΔB ≈ (V·D·T)/(N·Ae). Substituting into Steinmetz shows that raising f reduces ΔB per cycle (∝1/f) but f itself enters Pv as f^α (α typically >1), so the two do not simply cancel — this is the physical root of "higher frequency is not always cheaper" (the same thesis as our LLC frequency–knee article).

3.3 How loss maps to temperature rise

Loss ultimately shows up as temperature rise: ΔT ≈ P_total × R_th (thermal resistance). For the same loss, an open bench and a sealed potted enclosure differ enormously in thermal resistance — the physical reason behind the "passes on the bench, fails in the unit" trap discussed in Chapter 7.

DC Bias & Saturation: FeSi's Core Weakness

The most underrated property of FeSi powder cores is that permeability falls monotonically with DC field strength Hdc — an inherent "soft saturation" of distributed-gap powder cores, not a defect. The problem: many engineers size the inductor on zero–bias permeability (μi), ignoring that μ may have rolled off 30—60% at peak current.

4.1 L@Ipk retention is the real design anchor

For a typical μ60 FeSi core, permeability rolls off smoothly with Hdc (soft saturation), unlike the steep cliff of ferrite. This gives an engineering advantage — under overcurrent, inductance falls gradually rather than collapsing, so the system is harder to destabilize. But it also sets a trap: **if you size on zero-bias inductance, the actual inductance at peak current may be too low to hold the design ripple, pushing ripple current — and both core and winding loss — out of spec.**

DC bias Hdc (ref.)	Typical μ60 retention (ref.)	Engineering meaning
Low (~20 Oe)	≈ 90–95%	Near zero-bias, ample margin
Medium (~50 Oe)	≈ 70–80%	Compute L@rated current here
High (~100 Oe)	≈ 45–60%	Often reached in three-phase; check carefully
Heavy (>120 Oe)	< 45%	FeSi near its limit; evaluate a material change

The table is a public reference trend for typical μ60 FeSi; actual retention depends on grade, magnetic path and temperature — confirm per project with measured curves.

Why does the powder core soft-saturate? FeSi powder cores are made by coating FeSi alloy powder with an insulating binder and pressing it, leaving countless tiny gaps distributed throughout the body (distributed gap). These gaps make the effective reluctance rise gently with applied field, so permeability declines continuously with Hdc rather than collapsing at one point. Understand this and you understand both sides: soft saturation gives overcurrent tolerance, but it also means **the rated operating point must be evaluated on the already-reduced μ, never on initial μi.**

The correct design flow is reversed: fix peak current Ipk → compute Hdc = N·Ipk/le (le = magnetic path length) → read the μ retention at that point from the grade's measured μ–Hdc curve → back–calculate the turns N needed to meet L@Ipk using the reduced μ. **If you compute turns from μi, the inductance "shrinks" at peak current, ripple grows, and core and winding loss rise together — the common root of countless "simulation fine, measurement overheats" cases.**

4.2 Three-phase high-current material card: when to switch

When a three–phase 22 kW scenario drives DC bias high, FeSi's μ roll–off forces you to "add volume to keep inductance," losing the power–density advantage. Evaluate materials with higher Bsat and better bias retention. Public characteristics compared:

Material	FeSi	High Flux	Sendust
Saturation Bs (ref.)	≈ 1.5–1.6 T	≈ 1.5 T	≈ 1.0–1.05 T
DC bias retention	Good	Best (best at equal volume)	Medium
Core loss	Medium	Higher	Lowest
Relative cost	Low	High	Medium
Best fit	Single-phase / medium bias	Three-phase heavy bias, volume-critical	Low loss, medium bias need

Switch criterion (engineering rule of thumb, confirm per project): when the μ retention at the peak operating point (from measured curves) drops below ~60%, and holding L forces a significant increase in core volume or pushes rise out of spec, bring High Flux into the comparison. If the scenario is extremely core–loss–sensitive with only medium bias, Sendust's low loss is more valuable. The three are not substitutes — they are trade–offs chosen by operating point.

Single- vs Three-Phase: Differentiated Design

Applying one inductor–design logic to both single– and three–phase is a common engineering reflex. Their constraint priorities are not the same:

Design dimension	11 kW single-phase totem-pole	22 kW three-phase Vienna
Main constraint	HF AC loss and thermal rise	DC-bias retention and saturation margin
Inductance need	Higher (single-phase ripple control)	Can be relaxed via interleaving
Per-leg current / bias	High current, medium-high bias	Lower current per leg but still heavy bias
Flat-wire turn strategy	Cut AC loss first, control layers	Balance window use + per-unit consistency (×3)
Material first choice	μ60 FeSi	FeSi; evaluate High Flux at heavy bias
Consistency requirement	Single unit	Three matched units; tighten spread

Quantify "consistency" from slogan to controllable quantity. One root of three–phase current imbalance is inductance spread among the three units at the same bias. With L1, L2, L3 under the same drive voltage, the leg with the smallest inductance carries larger ripple current, higher average loss, higher rise; the higher rise then shifts that leg's core and winding operating point further — a positive feedback. In other words, **a few-percent inductance spread can amplify into a double-digit-percent per-leg loss difference and a 10°C-plus temperature gap.**

Three layers of control: (1) core — match by batch, control initial μ spread; (2) winding — edge–wound flat–wire tooling guarantees consistent turns and window position, far better than the random tension of hand–wound round wire; (3) sizing — leave enough post–bias margin so all three legs sit in the gentle region of the μ curve, not the steep zone. **This is the engineering substance of "three-phase is not single-phase ×3": the single-phase design answers for one unit; the three-phase design must answer for the spread.**

Two Complete Worked Examples

Important: all numbers in this chapter are demonstration values derived from public industry reference data — not ProMagTech measured data, and not representative of any specific product. Cores are described by generic specification, not vendor part numbers. Values illustrate design method only; real projects must be confirmed against the measured μ-Hdc curve and measured AL of the chosen grade and the customer's conditions.

Example A: 11 kW Single-Phase Totem-Pole PFC Inductor

Given conditions (demo):

Item	Demo value
Topology / power	Single-phase totem-pole PFC / 11 kW
Input voltage	230 Vac
Inductor RMS current Irms	≈ 47.8 A (≈ P/Vac)
Peak current Ipk	≈ 68 A (incl. ripple)
Target inductance L	≈ 40 µH
Switching frequency f	100 kHz
Candidate core	µ60 FeSi, large toroid (OD ~78 mm class, le≈18.4 cm, AL≈89 nH/N², public ref.)
Winding	Flat wire, section ≈1×6 mm (≈6 mm²)

Step 1 (wrong start): turns from zero-bias AL: $N = \sqrt{(L/AL)} = \sqrt{(40\mu H / 89nH)} \approx$ **21 turns (paper)**. But this is the zero-bias value.

Step 2: peak field strength $Hdc = N \cdot Ipk/le = 21 \times 68/0.184 \approx 7{,}760$ A/m ≈ **98 Oe** — in the clear roll-off region of µ60 FeSi.

Step 3: from the reference µ–Hdc trend, retention near 98 Oe is about 60%, so $\mu eff \approx 0.6 \cdot \mu i$.

Step 4 (recompute): $L \propto \mu \cdot N^2$, so holding 40 µH needs $N\_real \approx 21/\sqrt{0.6} \approx$ **27 turns**. Recompute $Hdc = 27 \times 68/0.184 \approx 125$ Oe, retention re-checked at ~55–58%; converges around 27–28 turns after one iteration.

Step 5 (loss check, demo): with 27 turns, MLT≈0.13 m, flat wire 6 mm²:

- **DC resistance DCR** = $\rho \cdot N \cdot MLT/A\_cu \approx 1.72e\text{-}8 \times 27 \times 0.13/6e\text{-}6 \approx$ **10 mΩ**
- **Winding DC copper loss** $P\_cu = Irms^2 \cdot DCR \approx 47.8^2 \times 0.010 \approx \approx$ **23 W** (DC component only; HF added loss is separate — flat wire lowering Rac/Rdc is key here)
- **Core loss** from ripple swing ΔB via Steinmetz, computed against the chosen grade's loss curve; logged as a to-be-measured item at the demo stage.

Example A conclusion: µ60 FeSi is viable for single-phase 11 kW — Hdc lands in the medium-high bias band but still inside the material's usable range; 27 turns and ~10 mΩ DCR is a reasonable engineering starting point. **Note: 23 W of DC copper loss is high, marking this as a "high-current, control-AC-loss-and-rise-with-flat-wire" scenario — exactly as Chapter 3 argued.**

Example B: 22 kW Three-Phase Vienna PFC Inductor (single leg)

Given conditions (demo):

Item	Demo value
Topology / power	Three-phase Vienna PFC / 22 kW (≈7.3 kW per phase)
Phase voltage	230 Vac (400 V line system)
Per-phase RMS current Irms	≈ 31.9 A
Peak current Ipk	≈ 45 A
Target inductance L	≈ 120 µH
Switching frequency f	100 kHz
Candidate core	µ60 FeSi, mid toroid (OD ~63 mm class, le≈14.3 cm, AL≈135 nH/N², public ref.)
Winding	Flat wire, section ≈1×4 mm (≈4 mm²), three matched units

Step 1: N = √(120μH/135nH) ≈ **30 turns (paper)**.

Step 2: Hdc = 30×45/0.143 ≈ 9,440 A/m ≈ **119 Oe** — already heavy bias.

Step 3: retention at 119 Oe is about 55%, μeff ≈ 0.55·μi.

Step 4: N_real ≈ 30/√0.55 ≈ **40 turns**. But recomputing Hdc = 40×45/0.143 ≈ **158 Oe, retention drops below ~50%** — a vicious circle: more turns → higher Hdc → lower μ → still more turns.

Step 5 (triggers the material-switch criterion): this is the real-world trigger point of the Chapter-4 criterion.

- **Forcing μ60 FeSi:** you keep piling turns / core volume, DCR rises to ≈17 mΩ, copper loss ≈17.5 W, and both volume and rise worsen — losing power density.
- **Correct move:** bring **High Flux** into comparison — its better DC-bias retention at equal volume keeps effective μ markedly higher than μ60 near 158 Oe, holding 120 μH with fewer turns and breaking the deadlock.

Example B conclusion: same method, higher power, opposite conclusion — **three-phase 22 kW pushes μ60 FeSi past its comfort zone.** Carry over Example A's "FeSi is enough" lesson and you hit a wall in three-phase. **This is the most concrete numerical evidence for "three-phase ≠ single-phase ×3."**

What the two examples teach side by side

Dimension	Ex. A single-phase 11 kW	Ex. B three-phase 22 kW / leg
Peak Hdc (real turns)	≈125 Oe	≈158 Oe
μ retention	≈55–58%	below 50%
FeSi conclusion	Usable (medium-high bias)	Past the limit; evaluate High Flux
Dominant loss tension	HF AC loss + thermal rise	μ roll-off + volume/rise spiral
Design action	Flat wire to control AC loss	Switch material or grow volume

Reminder: AL, retention, DCR and copper loss above are computed from public reference values to demonstrate method. Before production, re-verify against measured curves and measured AL, and validate temperature rise in the target thermal environment.

Power-Density Wall & Thermal Validation

The ceiling on power density is essentially "how much loss can you dissipate within the allowed temperature rise." OBCs are usually sealed / potted, with thermal resistance far higher than an open bench — so the same loss produces a higher rise in the unit.

- **Bench vs unit:** an open bench has good natural convection; a sealed enclosure relies on potting conduction and case transfer, with thermal resistance several times the bench.
- **Potting CTE mismatch & thermal cycling:** mismatch in CTE between potting compound and winding/core builds stress under thermal cycling, which over time can cause microcracks and degraded conduction. ProMagTech covers this in a dedicated article — see Chapter 9 resources.
- **Real power-density constraint:** it is not how small the core can be, but whether the heat path can carry the loss out. The effective levers are lowering AC loss (flat wire) and improving potting conduction, not merely shrinking the core.

A pragmatic validation principle: temperature data only means something in the target thermal environment. Run full-load rise and thermal-cycling validation with the same potting compound, enclosure and mounting as production, and record the hot-spot location (usually inner-layer winding or core center). Releasing on open-bench data alone is the institutional source of "bench pass, unit fail."

Parameter Traps & Failure Modes

Typical trap	Consequence & countermeasure
Sizing on DCR only	Low DCR ≠ low loss. At 65–140 kHz, added AC loss can dominate rise; compute Rac/Rdc and core loss.
Using zero-bias μ for L	Ignores DC-bias roll-off; actual L at peak is short, ripple over-spec, loss cascades.
Ignoring ripple swing ΔB	Core loss is driven by ΔB, not peak B; a wrong ripple estimate overheats the core.
Bench pass, unit fail	Sealed/potted thermal resistance doubles vs bench; validate in the target thermal environment.

Typical trap	Consequence & countermeasure
Copying single to three-phase	Ignores per-unit consistency and bias differences; causes leg imbalance and local hot spots.
One material for all	Forcing FeSi at heavy bias bloats volume and loses density; evaluate High Flux.

FAQ

Q: What is the first parameter to confirm when sizing an OBC PFC inductor?

A: Inductance retention at peak current (L@Ipk) and temperature rise in the target thermal environment — not zero-bias inductance or DCR.

Q: Is flat wire always better than round wire at 65–140 kHz?

A: On PFC inductors with high current density, many layers and tight thermal budgets, flat wire's AC-loss and thermal advantages are clear; on small low-current inductors the benefit may be modest. Evaluate by operating point.

Q: When is FeSi powder core no longer enough?

A: When μ retention at the peak point drops below ~60% and holding inductance forces a large volume increase or out-of-spec rise, evaluate higher-Bsat, better-bias materials such as High Flux.

Q: Can a three-phase design just be single-phase ×3?

A: No. You must additionally control the inductance consistency and bias differences of the three units, or you get phase-current imbalance and local overheating.

Q: Why does a sample pass the bench but overheat in the unit?

A: A sealed/potted enclosure has far higher thermal resistance than an open bench, so the same loss gives a higher rise. Validate in the target thermal environment and improve potting conduction.

Related Resources & Contact

Further reading (ProMagTech technical library):

- LLC/DCX resonant transformer “frequency knee”: why volume gains diminish above 400 kHz
- Potting CTE mismatch & thermal-cycling failure
- 800 V HVDC bus inductor design guide

Talk to the ProMagTech engineering team

If you are sizing a flat-wire inductor for an OBC PFC stage, send us your operating conditions — input/output voltage, phase count and power, switching frequency, peak current, target temperature rise and packaging method — and we will run a selection evaluation and custom design for your case.

SHENZHEN PROMAGTECH CO.,LTD. · 深圳市谱磁科技有限公司

Web: www.promagtech.com | Email: zyong@promagtech.cn | Tel: +86 13537658938