

FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

Design and Optimisation of Magnetic Couplers in Wireless Power Transfer Systems

Lourenço Ramos Moreira Pinto de Carvalho



Mestrado em Engenharia Eletrotécnica e de Computadores

Supervisor: Prof. Justino Miguel Ferreira Rodrigues

July 27, 2026

Resumo

A adoção generalizada de veículos elétricos e de sistemas industriais automatizados tem intensificado a procura por soluções de transferência de energia sem fios eficientes e fiáveis. No centro destes sistemas encontra-se o acoplador magnético, cuja conceção determina de forma crítica a eficiência global e a robustez operacional. Contudo, as práticas de projeto atuais permanecem frequentemente fragmentadas: baseiam-se em avaliações pontuais que não contemplam a variabilidade espacial inerente à utilização real, nomeadamente os desalinhamentos laterais e as variações do entreferro. Acresce que, embora as normas vigentes estabeleçam objetivos de desempenho, não prescrevem um procedimento sistemático para traduzir esses requisitos numa geometria física otimizada.

Esta dissertação propõe uma metodologia abrangente e estruturada em múltiplas etapas para o projeto e otimização de acopladores magnéticos em sistemas de transferência de energia indutiva. A abordagem integra modelação analítica, optimização numérica e simulação computacional de alta fidelidade num fluxo de trabalho unificado. Um aspeto central da metodologia é o tratamento explícito de toda a envolvente operacional tridimensional, garantindo que a robustez espacial é incorporada no projeto desde o início. A estrutura é intencionalmente flexível, acomodando diversas topologias de bobinas e escolhas de materiais, respeitando simultaneamente restrições elétricas e térmicas práticas.

A eficácia da metodologia proposta é validada através da construção e ensaio experimental de protótipos físicos. As medições realizadas em condições de alinhamento representativas confirmam a consistência do processo de optimização e demonstram a praticabilidade da abordagem.

A metodologia de projeto resultante oferece um percurso sistemático e compatível com as normas para o desenvolvimento de acopladores magnéticos de alto desempenho, colmatando uma lacuna crítica entre os requisitos teóricos de desempenho e a implementação física, num conjunto variado de aplicações de carregamento sem fios estacionário.

Palavras-chave: Transferência de Energia sem Fios, Acoplamento Indutivo, Projeto de Bobinas, Estrutura de Optimização, Análise de Desalinhamento, Análise por Elementos Finitos.

Abstract

The widespread adoption of electric vehicles and automated industrial systems has increased the demand for efficient and reliable Wireless Power Transfer (WPT) solutions. At the core of these systems lies the magnetic coupler, whose design critically determines overall efficiency and operational robustness. However, current design practices often remain fragmented, relying on single-point evaluations that fail to account for the spatial variability inherent in real-world use, such as misalignment and air-gap variations. Furthermore, while existing standards establish performance targets, they do not prescribe a systematic procedure to translate these requirements into an optimal physical geometry.

This dissertation proposes a comprehensive multi-stage methodology for the design and optimisation of magnetic couplers in Inductive Power Transfer (IPT) systems. The approach integrates analytical modelling, numerical optimisation, and high-fidelity computational simulation within a unified workflow. A key aspect of the methodology is its explicit treatment of the full three-dimensional operational envelope, ensuring that spatial robustness is embedded in the design from the outset. The framework is intentionally flexible, accommodating various coil topologies and material choices while respecting practical electrical and thermal constraints.

The effectiveness of the proposed pipeline is validated through the construction and experimental testing of physical prototypes. Measurements acquired under representative alignment conditions confirm the consistency of the optimisation process and demonstrate the practicality of the approach.

The resulting design methodology offers a systematic, standards-compatible pathway for the development of high-performance magnetic couplers, addressing a critical gap between theoretical performance requirements and physical implementation across a range of stationary wireless charging applications.

Keywords: Inductive Power Transfer, Inductive Coupling, Coil Design, Optimisation Framework, Misalignment Analysis, Finite Element Analysis.

UN Sustainable Development Goals

The United Nations Sustainable Development Goals (SDGs) [48] provide a global framework to achieve a better and more sustainable future for all. It includes 17 goals to address the world's most pressing challenges, including poverty, inequality, climate change, environmental degradation, peace, and justice.

This dissertation contributes to the global transition toward sustainable mobility and industrial automation by proposing a systematic approach to enhancing the performance of contactless energy delivery systems. By improving the efficiency and spatial flexibility of the magnetic interface, this work supports the broader adoption of electrified transport and autonomous systems, directly reducing the global reliance on carbon-intensive energy sources and manual intervention.

The specific Sustainable Development Goals addressed in this work are:

SDG 7 Ensure access to affordable, reliable, sustainable and modern energy for all.

SDG 9 Build resilient infrastructure, promote inclusive and sustainable industrialization and foster innovation.

SDG 13 Take urgent action to combat climate change and its impacts.

SDG	Target	Contribution	Performance Indicators and Metrics
7	7.3	Maximising the efficiency of the energy transfer interface reduces overall system losses and optimises energy consumption.	Target energy transfer efficiency significantly exceeding 90% in ideal conditions.
	7.a	Enhancing the spatial flexibility of wireless systems facilitates the global implementation of clean energy infrastructure.	Robustness of performance across a wide range of operational positions and distances.
9	9.4	Providing systematic design workflows enables industries to upgrade infrastructure for electric and autonomous vehicles.	Scalability of the design approach to different power levels and industrial applications.
	9.5	Advancing scientific research in magnetic coupling enables the development of high-performance components with optimised material usage.	Balanced design selection considering both efficiency gains and material resource requirements.
13	13.2	Accelerating the transition to zero-emission mobility by improving the convenience and reliability of charging processes.	Successful validation of system performance under realistic misalignment scenarios.

Acknowledgements

I am deeply grateful to Professor Justino Rodrigues, my supervisor, for his unwavering support, thoughtful guidance, and constant availability throughout this dissertation. His straightforward and pragmatic counsel proved invaluable, helping me navigate critical decisions with clarity and confidence.

I am equally thankful to Pedro Pascoal, my supervisor at INESC TEC, whose technical expertise and continuous encouragement were fundamental to this work. His willingness to engage in deep discussions and his perceptive feedback significantly shaped the direction of this research.

I am also grateful to all members of the SGEVL laboratory for fostering a stimulating and welcoming environment. The camaraderie and mutual support among colleagues made every day more enjoyable and productive.

A special word of appreciation goes to Alexandre Gomes and Ahmed Hussein for the countless enriching conversations and collaborative spirit we shared. Their insights, constructive criticism, and willingness to explore different perspectives greatly enhanced the quality of this dissertation.

To my family, I owe a debt of gratitude that words cannot fully express. Their unconditional belief in me, their patience, and their constant encouragement have been my anchor throughout this journey. This achievement belongs to them as much as it does to me.

Lastly, I want to thank my friends for their understanding, for the moments of laughter and distraction, and for reminding me that there is life beyond the thesis. Their presence kept me grounded and balanced when I needed it most.

Lourenço Carvalho

Contents

1	Introduction	1
1.1	Context and Motivation	1
1.2	Research Environment	3
1.3	Problem Statement	4
1.4	Objectives and Scope	5
1.5	Main Contributions	7
1.6	Thesis Outline	7
2	Fundamentals of Inductive Power Transfer Systems	9
2.1	System Architecture of an IPT Charger	9
2.2	Magnetic Induction and Coupling	10
2.3	Electrical Properties: Resistance, Skin and Proximity Effects	12
2.3.1	Skin Effect	12
2.3.2	Proximity Effect	13
2.3.3	Mitigation with Litz Wire	14
2.4	Electrical Properties: Self and Mutual Inductances	14
2.4.1	Self-Inductance	15
2.4.2	Mutual Inductance	15
2.4.3	Experimental Determination of Inductance	16
2.5	Metrics for Coupling Performance	17
2.5.1	Quality Factor	18
2.5.2	Coupling Coefficient	18
2.5.3	Implications for System Performance	18
2.6	Resonant Operation and Compensation Networks	19
2.6.1	Resonant Condition	20
2.6.2	Basic Compensation Topologies	20
2.6.3	Deductions for the SS Topology	21
2.6.4	Practical Aspects of Compensation	23
2.7	Ferrite Materials for Flux Guidance	23
2.7.1	Magnetic Properties of Ferrites	23
2.7.2	Common Ferrite Shapes for IPT Coils	24
2.7.3	Effect of Ferrite on Inductance and Coupling	25
2.8	Electromagnetic Shielding	26
2.8.1	Passive Shielding	26
2.8.2	Active Shielding	27
2.8.3	Reactive Shielding	27
2.9	Finite Element Analysis and Thermal Simulation Tools	28
2.9.1	Available FEA Software for IPT Design	28

2.9.2	Anslys Maxwell for Electromagnetic Simulation	28
2.9.3	Anslys Icepak for Thermal Simulation	29
2.10	Summary	30
3	State-of-the-Art and Literature Review	31
3.1	Overview of IPT Systems	31
3.2	Coil Geometries and Magnetic Coupler Topologies	32
3.2.1	Circular Coils	32
3.2.2	Rectangular and Square Coils	33
3.2.3	Double-D (DD) and Related Topologies	34
3.2.4	Hybrid and Multi-Coil Designs	36
3.3	Ferrite Structures and Electromagnetic Shielding	36
3.3.1	Ferrite Arrangements: From Heuristics to Optimisation	36
3.3.2	Shielding Techniques and Their Trade-Offs	38
3.3.3	Critical Synthesis: Why Co-Design Matters	38
3.4	Misalignment and Coupling Variability	39
3.4.1	Coupling Degradation Under Misalignment	39
3.4.2	Implications for Efficiency and Control	40
3.5	SAE J2954 Standard: Critical Analysis	41
3.5.1	What the Standard Provides	41
3.5.2	Critical Analysis: What the Standard Does Not Specify	42
3.5.3	How the Present Work Uses the Standard	45
3.6	Design Optimisation Approaches	45
3.6.1	Finite Element Analysis: Accurate but Computationally Intensive	45
3.6.2	Analytical Formulas: Fast but Limited in Scope	46
3.6.3	Heuristic and Rule-of-Thumb Designs	46
3.6.4	Surrogate-Assisted and Multi-Objective Optimisation	47
3.7	Positioning of the Present Work	47
4	Analytical Toolchain and Validation	50
4.1	Rationale for the Toolchain	50
4.2	Tool 1 – Litz Wire Dimensioning	50
4.2.1	Mathematical Foundation	51
4.2.2	Typical Use and Example	52
4.3	Tool 2 – Single Coil Analysis	52
4.3.1	Geometry Discretisation and Conductor Length	53
4.3.2	DC and AC Resistance	53
4.3.3	Self-Inductance and Quality Factor	54
4.3.4	Typical Use and Example	54
4.4	Tool 3 – Spatial Coupling Analysis	55
4.4.1	Spatial Sweep and Coupling Metrics	55
4.4.2	Exploitation of Symmetry and Grid Discretisation	55
4.4.3	Statistics and Performance Optimisation	57
4.4.4	Typical Use and Example	58
4.5	Tool 4 – Coupling Optimisation	59
4.5.1	Why Optimisation Is Needed	59
4.5.2	Particle Swarm Optimisation (PSO) – Overview	60
4.5.3	Parameter Space and Discretisation	61
4.5.4	Objective Function and Constraints	61

4.5.5	Typical Use and Example	62
4.6	Validation of Air-Core Predictions	63
4.6.1	Test Matrix Overview	63
4.6.2	Conductor Constructions	64
4.6.3	Evaluation Methodology	65
4.6.4	Comparison with Ansys Maxwell	65
4.6.5	Validation Closure	67
4.7	Critical Assessment and Future Directions	67
4.7.1	Ferrite and Shielding Are Excluded	68
4.7.2	No Formal Optimality Guarantee	69
4.7.3	Limited Multi-Receiver Optimisation	69
4.7.4	No Angular Misalignment	70
4.7.5	Finite Coil Geometry Support	70
4.7.6	Lack of Parameter Sensitivity Analysis	71
4.8	Chapter Summary	71
5	Complete Methodology and Optimisation Results	73
5.1	Five-Stage Design Pipeline: Overview	73
5.2	Design Constraints and Specifications	75
5.2.1	System-Wide Specifications	75
5.2.2	Misalignment Envelope and Evaluation Grid	75
5.2.3	Remaining Design Inputs	75
5.3	Stage 1: MATLAB Toolchain Pre-Filtering	76
5.3.1	Litz Wire Dimensioning (Tool 1)	76
5.3.2	Coupling Optimisation (Tool 4)	77
5.3.3	Optimisation Results	78
5.3.4	Proximity Loss Diagnosis	79
5.4	Stage 2: FEA Air-Core Validation	81
5.5	Stage 3: Ferrite and Shielding Integration	82
5.5.1	Design Standardisation Strategy	82
5.5.2	Ferrite Tile Selection and Material Properties	83
5.5.3	Aluminium Shield Specification	83
5.5.4	Standardised Ferrite and Shield Parameters	84
5.5.5	Coupling $A-I$ as a Non-Ideal Baseline	85
5.5.6	Per-Coupling Ferrite Layout Definition	85
5.5.7	Spatial Layout Selection Methodology and Cost-Performance Metric	87
5.5.8	Effect of Ferrite and Shield on Electrical Parameters	88
5.6	Stage 4: Electrical Worst-Case Assessment	89
5.6.1	Worst-Case Electrical Analysis	90
5.6.2	Ferrite Saturation Verification	91
5.6.3	Resonance Detuning Sensitivity	92
5.7	Stage 5: Thermal Simulation	93
5.8	Design Closure	94
5.8.1	Litz Wire Electrical and Thermal Validation	95
5.8.2	Inductance Limit Validation	95
5.8.3	Final Coupler Cost Summary	96
5.9	Critical Assessment and Future Directions	96
5.9.1	Standardised Ferrite Layouts	97
5.9.2	Automation and Scalability Limitations	98

5.9.3	Thermal Simulation Assumptions	98
5.9.4	SS Topology Dependency	99
5.9.5	Absence of Tolerance Analysis	99
5.9.6	Absence of SAE J2954 Reference Geometry Comparison	100
5.10	Chapter Summary	100
6	Practical Implementation and Validation	102
6.1	Introduction	102
6.2	Mechanical Test Jig Design	102
6.2.1	Overall Structure and Three-Axis Movement	103
6.2.2	Tx and Rx Sandwich Assembly	104
6.2.3	Thermal Considerations	104
6.3	Coil Holders and Modular Winding	105
6.4	Soldering and Crimping of Litz Wire	107
6.5	Assembled Coils and Test Jig	108
6.6	Measurement of DC Resistance	109
6.6.1	Two-Wire versus Four-Wire Method	109
6.6.2	Measurement Results	110
6.6.3	Discussion of Discrepancies	110
6.7	Measurements of Individual Coils	111
6.7.1	Measurement Approach	111
6.7.2	Measurements of Air-Core Coils	112
6.7.3	Measurements with Ferrite and Aluminium Shield	113
6.8	Full Coupling Measurements	114
6.8.1	Measurement Approach	114
6.8.2	Measured Mutual Inductance	114
6.8.3	Extrapolated Coupling Coefficient and Maximum Efficiency	115
6.9	Critical Assessment and Future Directions	115
6.9.1	Measurement-Induced Errors	116
6.9.2	Ferrite-Related Uncertainties	116
6.9.3	Limitations of the Test Jig	117
6.9.4	Scope Limitations	117
6.10	Chapter Summary	118
7	Conclusions and Future Work	119
7.1	The Analytical Toolchain	119
7.2	The Complete Design Methodology	120
7.3	Experimental Implementation	121
7.4	Publication and Dissemination	121
7.5	Final Closing Statement	122
	References	124
A	Validation Couplings Configurations	129
A.1	T1: Single-layer Spiral – Single-layer Spiral	129
A.2	T2: Single-layer Spiral – Single-layer Spiral	130
A.3	T3: Single-layer Spiral – Single-layer Spiral	130
A.4	T4: Single-layer DD – Single-layer Spiral	130
A.5	T5: Single-layer DD – Single-layer DD	131

A.6	T6: Multi-layer Spiral – Single-layer Spiral	131
A.7	T7: Multi-layer Spiral – Multi-layer Spiral	131
A.8	T8: Multi-layer DD – Single-layer Spiral	131
A.9	T9: Multi-layer DD – Multi-layer DD	132
A.10	T10: Mixed conductor and topology pairs	132
B	Ferrite Layouts and Associated Metrics	133
B.1	Coupling <i>A-I</i>	133
B.2	Coupling <i>A-II</i>	134
B.3	Coupling <i>B</i>	134
B.4	Coupling <i>C</i>	135
B.5	Coupling <i>D</i>	135

List of Figures

1.1	Classification of wireless power transfer technologies. Adapted from [35].	2
1.2	Comparison of charging methods: (a) conventional wired charging, (b) wireless IPT for EV parking, (c) wireless IPT for AGVs in industrial settings.	3
1.3	Types of misalignment in static IPT charging: (a) lateral displacement, (b) vertical displacement, (c) angular tilt.	4
2.1	Block diagram of a generic IPT system. Adapted from [9].	10
2.2	Magnetic flux distribution in a loosely coupled IPT system. Adapted from [5]. . .	11
2.3	Illustration of current distribution in a conductor: (a) DC, (b) skin effect at high frequency, (c) both skin and proximity effects combined. The colour scale ranges from low (blue) to high (red) current density. Adapted from [13].	13
2.4	Structure of a Litz wire: (a) 3-D model, (b) real Litz wire with visible bundle contours. Adapted from [13].	14
2.5	Discretisation of a planar spiral coil into small straight segments for numerical evaluation: (a) overall coil geometry, (b) discretised segments.	17
2.6	Measuring mutual inductance using the series-aiding (a) and series-opposing (b) method.	18
2.7	Basic compensation topologies: (a) Series-Series (SS), (b) Series-Parallel (SP), (c) Parallel-Series (PS), (d) Parallel-Parallel (PP).	20
2.8	Equivalent circuit of an SS-compensated IPT system.	21
2.9	Nonlinear B-H characteristic of a typical ferrite material. Adapted from [14]. . .	24
2.10	Common ferrite structures: (a) planar solid tile, (b) segmented I-bars, (c) U-core, (d) E-core.	25
2.11	Classification of electromagnetic shielding techniques for IPT systems: (a) passive shielding, (b) active shielding, (c) reactive shielding. Adapted from [25].	26
2.12	Example comparison of (a) the initial tetrahedral mesh and (b) the adaptively refined mesh after several passes. Adapted from [5].	29
3.1	Common IPT coil topologies (from left to right): Double-D Quadrature (DDQ), Double-D (DD), rectangular and circular. Adapted from [61].	32
3.2	Example of a circular coil geometry (a) and the corresponding magnetic field vector distribution in the YZ plane (b).	33
3.3	Example of a rectangular coil geometry (a) and the corresponding magnetic field vector distribution in the YZ plane (b).	34
3.4	Example of a DD coil geometry (a) and the corresponding magnetic field vector distribution in the YZ plane (b).	35
3.5	Comparison of coil performance under misalignment conditions: (a) lateral misalignment and (b) vertical misalignment. Adapted from [58].	40

3.6	Coordinate axes for SAE J2954 misalignment definitions. Adapted from [43] and [38].	42
3.7	Example coil geometries from SAE J2954 appendixes: (a) universal GA (for all power and Z-classes) and (b) test station VA for WPT1/Z1 class. Adapted from [43].	44
4.1	3D views of the example spiral coil: (a) angled view, (b) top view.	55
4.2	Point cloud of receiver coil centres (red) for a spatial analysis: (a) full region (189 points); (b) first quadrant (60 points). The transmitter coil centre is fixed at the origin (blue).	57
4.3	3D visualisation of Tool 3's Tx (blue) and Rx (red) coils at a given offset (50 mm lateral in x and y, 150 mm air gap).	59
4.4	Illustration of particle swarm evolution: (a) initial random positions, (b) convergence starting after several iterations.	60
4.5	3D visualisation of Tool 4's Tx (blue) and Rx (red) coils at a given offset (50 mm lateral in x and y, 150 mm air gap).	64
4.6	Boxplots of percent errors for all direct parameters (R , L , M) across the entire validation dataset.	68
5.1	Three-dimensional coil visualisations from the MATLAB toolchain: (a) <i>A-I</i> , (b) <i>A-II</i> , (c) <i>B</i> , (d) <i>C</i> , (e) <i>D</i>	80
5.2	Magnetic field distribution (side view, YZ -plane) in the air-core configurations: (a) symmetric coupling <i>A-II</i> ; (b) asymmetric coupling <i>B</i>	82
5.3	Coil-to-coil efficiency maps over the evaluation grid of the air-core configurations: (a) symmetric coupling <i>A-II</i> ; (b) asymmetric coupling <i>B</i>	82
5.4	Commercial Mn-Zn ferrite tiles geometries selected for coupler construction: (a) 100×100 mm, (b) 150×100 mm.	83
5.5	Relative permeability versus frequency for the selected Mn-Zn ferrite (typical curve from supplier datasheet).	84
5.6	Ferrite tile layout for the coupler of coupling <i>A-I</i> (<i>Type X</i>).	85
5.7	Ferrite tile layouts for 500×500 mm coils: (a) <i>Type I</i> , (b) <i>Type II</i> , (c) <i>Type III</i> , (d) <i>Type IV</i> , (e) <i>Type V</i> , (f) <i>Type VI</i> , (g) <i>Type VII</i> , (h) <i>Type VIII</i>	86
5.8	Ferrite tile layouts for 350×350 mm receiver coils: (a) <i>Type I</i> , (b) <i>Type II</i> , (c) <i>Type III</i> , (d) <i>Type IV</i> , (e) <i>Type V</i> , (f) <i>Type VI</i>	86
5.9	Magnetic field distribution (side view) with <i>Type IV</i> ferrite and aluminium shield: (a) coupling <i>A-II</i> ; (b) coupling <i>B</i>	89
5.10	Typical frequency response of a series <i>LC</i> circuit near resonance (magnitude in red, phase in blue).	93
5.11	Steady-state temperature distribution for coupling <i>A-II</i> under worst-case losses: (a) Litz wire coils, (b) ferrite tiles, (c) aluminium shield.	93
5.12	Steady-state temperature distribution for coupling <i>B</i> under worst-case losses: (a) Litz wire coils, (b) ferrite tiles, (c) aluminium shield.	94
5.13	Photograph of the two purchased Litz wire constructions.	96
6.1	3D CAD model of the test jig. The arrows indicate the x , y and z directions of movement.	103
6.2	Exploded view of a sandwich (transmitter of coupling <i>A-II</i>).	104
6.3	CAD models of the three corner holders: (a) <i>A-I</i> , (b) <i>A-II</i> , (c) <i>B-Rx</i>	106
6.4	Winding steps for the <i>A-II</i> transmitter: (a) final stage of the winding process, (b) completed coil.	106

6.5	Comparison of soldered terminations: (a) poor, (b) almost perfect, (c) perfect. . .	107
6.6	Completed termination of a coil.	108
6.7	Top acrylic sheets with wound coils: (a) <i>A-I</i> , (b) <i>A-II</i> , (c) <i>B-Rx</i>	108
6.8	Physical implementations of: (a) assembled sandwich (transmitter of coupling <i>A-II</i>) and (b) test jig in the laboratory (receiver side aluminium shield removed for visibility).	109
6.9	Comparison of (a) two-wire and (b) four-wire Kelvin measurement methods. Adapted from [2].	110
6.10	Impedance measurement and model fitting for a representative coil. (a) Measured magnitude (red) and phase (blue). (b) Model B and the estimated parameters. . .	112
6.11	Measurement setup showing the reference planes (before and after calibration and after the manual zeroing step).	113

List of Tables

3.1	SAE J2954 wireless power transfer (WPT) power classes.	41
3.2	SAE J2954 Z-classes for ground clearance (VA to GA).	41
3.3	SAE J2954 misalignment tolerances.	41
3.4	SAE J2954 Class I minimum system efficiency requirements.	42
4.1	User inputs for Tool 1.	52
4.2	Outputs for Tool 1.	52
4.3	Example inputs for Tool 2.	56
4.4	Example outputs for Tool 2.	56
4.5	Number of evaluation points for the SAE J2954 Class <i>ZI</i> envelope (step size 25 mm).	56
4.6	Evaluation region for Tool 3 example.	58
4.7	Input parameters for Tx and Rx coils.	58
4.8	Coupling statistics for Tool 3 example.	59
4.9	Input coil search space for Tool4 example.	62
4.10	Optimised Rx parameters for Tool 4 example.	63
4.11	Coupling statistics for Tool 4 example.	63
4.12	Summary of test coupling configurations for validation.	63
4.13	Summary of Litz wire constructions used in validation test matrix.	64
4.14	Percent errors for resistance predictions, compared to Ansys Maxwell.	66
4.15	Percent errors for inductance predictions, compared to Ansys Maxwell.	67
4.16	Errors for coupling and efficiency predictions, compared to Ansys Maxwell.	67
5.1	Evaluation grid for spatial sweeps (first quadrant of SAE J2954 Z1 envelope).	75
5.2	Typical current density recommendations for copper windings in IPT systems.	76
5.3	Litz wire construction parameters from Tool 1 for both the coarse-strand (300×AWG32) and fine-strand (1200×AWG38) candidates.	77
5.4	Stage 1 optimised geometric parameters and air-core electrical metrics for all five coupler candidates.	79
5.5	Proximity factor comparison for the <i>A-I</i> geometry evaluated with both Litz wire constructions.	80
5.6	Stage 2 air-core electrical metrics for all five couplings from Ansys Maxwell across the evaluation grid.	81
5.7	Magnetic properties of the selected Mn-Zn ferrite material.	83
5.8	Standardised ferrite tile and aluminium shield assembly parameters applied uniformly across all five coupler designs.	84
5.9	Reference unit costs for ferrite tiles and aluminium shield material used in the layout cost-performance metric.	87

5.10	Selected ferrite layouts and cost-performance metrics for each coupling, spatially averaged over the evaluation grid.	88
5.11	Stage 3 electrical metrics with ferrite and aluminium shield.	89
5.12	Stage 4 worst-case electrical parameters for all five couplings under rated output of 3.7 kW.	91
5.13	Peak ferrite flux density $B_{\max}$ (mT) at the aligned minimum air-gap position for all five couplings.	92
5.14	Stage 5 maximum component hotspot temperatures ($^{\circ}\text{C}$) for all five couplings under worst-case power dissipation with natural convection cooling.	94
5.15	Validation of Litz wire electrical and thermal requirements against supplier specifications.	95
5.16	Bill-of-materials cost breakdown (€) for each coupler design.	97
6.1	Measured DC resistance compared with simulation results.	110
6.2	Measured total resistance and self-inductance for air-core coils.	113
6.3	Measured total resistance and self-inductance for fully assembled coils.	114
6.4	Measured mutual inductance compared with simulation results.	115
6.5	Relative and absolute errors of the coupling coefficient and maximum efficiency for each test coupling.	115
A.1	T1 configurations	129
A.2	T2 configurations	130
A.3	T3 configurations	130
A.4	T4 configurations	130
A.5	T5 configurations	131
A.6	T6 configurations	131
A.7	T7 configurations	131
A.8	T8 configurations	131
A.9	T9 configurations	132
A.10	T10 configurations	132
B.1	Ferrite analysis for Coupling <i>A-I</i>	133
B.2	Ferrite analysis for Coupling <i>A-II</i>	134
B.3	Ferrite analysis for Coupling <i>B</i>	134
B.4	Ferrite analysis for Coupling <i>C</i>	135
B.5	Ferrite analysis for Coupling <i>D</i>	135

List of Acronyms

AC	Alternating Current
AGV	Automated Guided Vehicles
AWG	American Wire Gauge
BNC	Bayonet Neill-Concelman (coaxial cable connector)
CAD	Computer-Aided Design
CFD	Computational Fluid Dynamics
CPES	Centre for Power and Energy Systems
DC	Direct Current
DD	Double-D (coil topology)
DDQ	Double-D Quadrature (coil configuration)
EMF	Electromotive Force
ESR	Equivalent Series Resistance
EV	Electric Vehicle
FEA	Finite Element Analysis
FDM	Fused Deposition Modeling
GA	Ground Assembly
ICNIRP	International Commission on Non-Ionizing Radiation Protection
INESC TEC	Institute for Systems and Computer Engineering, Technology and Science
IPT	Inductive Power Transfer
LCC	Inductor-Capacitor-Capacitor (compensation network topology)
LCR	Inductance, Capacitance, and Resistance (meter)
MEPT	Maximum Efficiency Point Tracking
PETG	Polyethylene Terephthalate Glycol
PI	Polyimide
PLA	Polylactic Acid
PP	Parallel-Parallel (compensation network topology)
PS	Parallel-Series (compensation network topology)
PSO	Particle Swarm Optimization
RMS	Root Mean Square
RMSRE	Root Mean Square Relative Error
Rx	Receiver (coil)
SAE	Society of Automotive Engineers
SGEVL	Smart Grids and Electric Vehicles Laboratory
SP	Series-Parallel (compensation network topology)
SS	Series-Series (compensation network topology)
Tx	Transmitter (coil)
VA	Vehicle Assembly
WPT	Wireless Power Transfer

Chapter 1

Introduction

1.1 Context and Motivation

The global transition towards sustainable mobility and the increasing automation of industrial and logistics operations have intensified the demand for efficient, reliable, and user-friendly energy delivery solutions. Electric Vehicles (EVs), Automated Guided Vehicles (AGVs), and autonomous mobile robots are central to this transformation, yet their widespread adoption remains constrained by the limitations of conventional wired charging. Physical connectors and cables present several drawbacks: they require manual intervention, are susceptible to wear and environmental contamination, pose safety hazards such as electric shock and sparking, and restrict operational flexibility in dynamic or autonomous environments [1, 28].

Wireless Power Transfer (WPT) has emerged as a compelling alternative, enabling contactless energy delivery across an air gap without any physical connection. Far from being a nascent technology, WPT is already embedded in everyday life at low power levels. The Qi standard – developed by the Wireless Power Consortium and now integrated into virtually all modern smartphones, earbuds, and smartwatches – uses near-field magnetic induction to deliver 5–15 W across gaps of a few millimetres, eliminating charging cables in consumer electronics [63]. Medical implants such as cochlear processors and cardiac pacemakers rely on transcutaneous inductive coupling to recharge internal batteries without piercing the skin, a safety-critical advantage in clinical settings [35]. These commercially proven deployments at low power demonstrate the practical viability and reliability of near-field magnetic coupling – the challenge addressed in this dissertation is scaling the same principle to the kilowatt levels required for automotive and industrial applications, where air gaps are orders of magnitude larger and efficiency requirements far more stringent.

WPT technologies are broadly classified into two families according to their operating mechanism and effective range, as illustrated in Fig. 1.1.

Radiative far-field WPT propagates electromagnetic energy through freely travelling waves and is suited to long-range applications (metres to kilometres). Two sub-technologies are distinguished: **Microwave Power Transfer (MPT)**, which uses microwave beams for applications such




Figure 1.1: Classification of wireless power transfer technologies. Adapted from [35].

as satellite-to-ground power beaming and long-range sensor supply, and **Laser Power Transfer (LPT)**, which uses directed laser beams for line-of-sight scenarios such as drone recharging. Because radiated power spreads over large solid angles and represents a potential exposure hazard, both MPT and LPT are impractical for the kilowatt-scale, close-proximity charging scenarios considered in this dissertation.

Non-radiative near-field WPT exploits the reactive electromagnetic field in the immediate vicinity of a source. At distances comparable to the device dimensions – typically millimetres to a few hundred millimetres – energy can be transferred with high efficiency and negligible far-field radiation. Near-field WPT subdivides into two technologies:

- **Capacitive Power Transfer (CPT):** energy is coupled through the electric field between two pairs of parallel conductive plates. CPT is compact and tolerates metallic obstacles between the plates, but the large voltages required to transfer significant power through air limit practical systems to low-power applications (typically below 1 kW) at very small gaps (below a few millimetres) [63].
- **Inductive Power Transfer (IPT):** energy is coupled through a shared magnetic field between a transmitter coil and a receiver coil. IPT can be implemented in non-resonant or resonant form. In non-resonant IPT, the coils are driven directly without compensation, which is simple but results in significant reactive power circulation and poor efficiency at the air gaps typical of EV charging. In resonant IPT, compensation capacitors are added to both the transmitter and receiver circuits, tuning the system to a specific operating frequency and cancelling the coil reactances at resonance; this dramatically reduces reactive power and maximises power transfer efficiency [16]. Resonant IPT is the standard approach for all medium- and high-power applications, including light-duty EV charging under the Society of Automotive Engineers (SAE) J2954 [43] standard. Throughout this dissertation, the term IPT refers exclusively to resonant implementations.

For stationary EV and AGV charging at power levels of 3–11 kW and air gaps of 100–250 mm, resonant IPT is the only technology that simultaneously satisfies the efficiency, power density, and safety requirements demanded by automotive and industrial standards. CPT cannot deliver the required power at these gaps. Resonant IPT therefore represents the natural and well-established choice, corroborated by the proliferation of commercial IPT charging systems and the publication of SAE J2954 as the governing standard for light-duty EVs [16, 17, 33]. Fig. 1.2 contrasts conventional wired charging with wireless IPT charging for both EV and AGV applications.



Figure 1.2: Comparison of charging methods: (a) conventional wired charging, (b) wireless IPT for EV parking, (c) wireless IPT for AGVs in industrial settings.

A typical IPT system for stationary charging comprises four functional blocks: a ground-side power electronic inverter that converts grid Alternating Current (AC) power to a high-frequency current; a transmitter (Tx) magnetic coupler consisting of a planar coil, ferrite flux guides, and electromagnetic shielding; a vehicle-side receiver (Rx) magnetic coupler of similar construction; and a rectifier and DC-DC converter that conditions the received power for the battery. The magnetic coupler pair – the transmitter and receiver pads together – is the passive, non-controllable heart of the system: once installed, its geometry and materials determine all performance outcomes that the active electronics must subsequently manage. It is this sub-system that is the primary focus of this dissertation.

The methodology developed in this dissertation, while demonstrated within the SAE J2954 framework for light-duty EVs, addresses this challenge in a general and systematic way that is directly applicable to any near-field IPT scenario where spatial robustness and high efficiency must be achieved simultaneously.

1.2 Research Environment

The research presented in this dissertation was developed at INESC TEC (Institute for Systems and Computer Engineering, Technology and Science) in Porto, Portugal, within the Centre for Power and Energy Systems (CPES). The work was primarily carried out in the Smart Grids and Electric Vehicles Laboratory (SGEVL), a facility equipped with specialised instrumentation for high-frequency power electronics.

This work represents a pioneering effort within INESC TEC, serving as a pilot project to establish a specialised design and optimisation competence for the magnetic components of IPT

systems. As a new research line for the centre, a primary goal of this dissertation is to validate the feasibility and scalability of the proposed design methodology, creating a robust baseline for more complex high-power and dynamic charging projects in the future.

It is important to emphasise that this dissertation is situated within a broader collaborative framework involving a multidisciplinary team at CPES. While this work focuses exclusively on the design, optimisation, and experimental validation of the magnetic coupler (the coils, ferrite guides, and shielding), other team members are concurrently developing the power electronic converters, resonant compensation networks, and active control strategies required for a complete WPT system. By isolating the magnetic interface as a standalone research objective, this pilot project provides the necessary technical foundation to ensure that future integrated systems at INESC TEC achieve the high efficiency and spatial robustness demanded by industrial and automotive standards.

1.3 Problem Statement

The magnetic coupler – comprising the transmitter and receiver coils, ferrite cores, and electromagnetic shielding – is the most critical subsystem in an IPT system. Its design directly determines power transfer efficiency, spatial robustness, system weight, material cost, and electromagnetic compatibility [11, 32, 55]. Unlike power electronic stages, which can be actively controlled, the magnetic coupling interface is passive and its performance is fundamentally dictated by geometric and material parameters.

In real-world stationary charging applications, perfect alignment between the transmitter and receiver coils is rarely achieved. Parking tolerances, vehicle suspension variations, and ground pad installation inaccuracies introduce lateral misalignments and air-gap deviations that significantly degrade coupling. The SAE J2954 standard for light-duty EV wireless charging defines vertical air gaps (Z-classes) ranging from 100 mm to 150 mm (Z1 class) and lateral misalignments of up to ± 75 mm in the vehicle length direction and ± 100 mm in the width direction [43]. Within this three-dimensional operational envelope, the coupling coefficient k can vary substantially, leading to efficiency degradation that may exceed 40% under unfavourable conditions, as documented in the literature [23, 24]. The three canonical misalignment modes are illustrated in Fig. 1.3.




Figure 1.3: Types of misalignment in static IPT charging: (a) lateral displacement, (b) vertical displacement, (c) angular tilt.

Despite extensive research on IPT couplers, several critical gaps remain in the existing literature. First, optimisation efforts are often fragmented, treating coil geometry, ferrite placement, and shielding as independent design variables rather than as coupled subsystems that must be co-optimised [10, 42]. Second, many studies evaluate coupler performance at a single nominal alignment position, neglecting the spatial variability that characterises practical operation [23, 24]. Third, while Finite Element Analysis (FEA) tools such as Ansys Maxwell and COMSOL Multiphysics provide high accuracy, their computational expense makes them unsuitable for exhaustive parametric exploration. Blind trial-and-error simulation approaches are time-consuming and do not guarantee convergence to optimal designs.

A fourth gap, equally significant, exists at the level of the governing standard itself. SAE J2954 specifies mechanical dimensions, electrical characteristics, and efficiency targets for reference coil geometries, providing a critical interoperability framework. However, it does not prescribe any systematic procedure for arriving at those geometries, nor does it define quantitative design objectives – such as a minimum spatial average of the coupling coefficient over the misalignment envelope – that a designer could use to guide optimisation. While standards provide normative reference designs, they do not prescribe a formal optimisation methodology, leaving a gap between performance requirements and the physical design process. The standard defines *what* a compliant coupler must achieve; it does not specify *how* to design one. This absence of a standardised design methodology means that each research group and manufacturer must develop their own optimisation approach from scratch, with no principled basis for comparing their results against those in the literature.

These gaps motivate the need for a systematic, multi-stage design methodology that combines analytical pre-design, constrained optimisation, and targeted high-fidelity simulation. Such an approach must explicitly account for the full three-dimensional operational envelope defined by applicable standards, enabling robust coupler designs that maintain high efficiency across realistic misalignment and air-gap variations.

1.4 Objectives and Scope

The primary objective of this dissertation is to develop, implement, and validate a complete multi-stage methodology for the design and optimisation of magnetic couplers in stationary inductive power transfer systems. This methodology integrates analytical modelling, numerical optimisation, finite element simulation, and experimental validation within a coherent workflow.

The specific objectives are as follows:

- Establish a system-level understanding of IPT operation, defining the boundary conditions and constraints imposed on the magnetic coupler by power electronics and compensation networks.

- Develop an analytical modelling framework capable of rapidly computing self-inductance, mutual inductance, coupling coefficient, and AC resistance (including skin and proximity effects) for planar spiral coils.
- Implement a constrained optimisation routine using Particle Swarm Optimisation (PSO) to maximise average coupling efficiency over the SAE J2954-defined misalignment region while respecting geometric and electrical constraints.
- Extend the framework to support multiple coil topologies, including circular, square, rectangular, and Double-D (DD) configurations.
- Validate optimised designs using high-fidelity three-dimensional finite element analysis (using Ansys Maxwell) and extract equivalent circuit models for system-level verification.
- Construct physical prototypes and perform experimental measurements to validate simulated performance and quantify discrepancies.

This work is part of a larger collaborative project. The design of power electronic converters, compensation networks, and active control strategies is being carried out by other team members. Consequently, the present thesis focuses exclusively on the magnetic coupler – its geometry, materials, and validation – assuming that a suitable Series-Series (SS) compensated inverter and rectifier exist. The SS topology is adopted throughout because its resonant frequency is uniquely independent of both load resistance and coupling coefficient variations, making it inherently robust for static EV charging where the coupling changes with every parking event. The coupler efficiency metric used throughout is derived under ideal resonant and impedance-matched conditions for the SS topology, which justifies this separation of concerns. Moreover, the methodology is not tied to a specific vehicle type: the constraints are taken from SAE J2954 (WPT1 power class, Z1 air gap 100–150 mm) as representative of many stationary IPT applications, including AGVs and industrial robots. No fixed outer dimensions for the coupler are assumed; the optimisation determines optimal sizes within practical limits.

The scope of this work is deliberately focused on static IPT systems with power levels up to 3.7 kVA, corresponding to SAE J2954 WPT1 class. The operational region follows SAE J2954 Z1 air gaps (100–150 mm) and lateral misalignment tolerances of ± 75 mm in x and ± 100 mm in y . Ferrite cores and aluminium electromagnetic shielding are included in the modelling and optimisation. Practical design constraints and a current density limit for passive cooling are also considered as part of the optimisation. Importantly, while SAE J2954 Z1 is used as the reference envelope throughout, the methodology itself is region-flexible: by redefining the evaluation grid, the same analytical toolchain and five-stage pipeline can be applied directly to any arbitrary three-dimensional operational envelope – for instance, the tighter alignment specifications of a docking robot, the wider tolerances of an industrial floor-mounted AGV charger, or future Z2/Z3 vehicle classes requiring air gaps up to 250 mm.

The following aspects are explicitly excluded from the scope: design of power electronic converters (rectifiers and inverters), implementation of active control strategies, detailed compensation network design (beyond establishing equivalent circuit models for efficiency evaluation), and dynamic or quasi-dynamic IPT systems. The optimised magnetic coupler is broadly applicable as the coil-to-coil efficiency depends mainly on the coupling coefficient and the quality factors of the coils, rather than on the specific compensation topology.

1.5 Main Contributions

This dissertation makes the following original contributions to the field of inductive power transfer systems:

- **Implementation of an analytical proximity-loss model:** A computationally efficient implementation of existing skin- and proximity-effect loss models for planar spiral coils, integrated into a constrained optimisation framework, enabling accurate loss estimation without recourse to time-consuming FEA.
- **Integrated MATLAB toolchain for coupler optimisation:** A flexible framework that implements geometry parameterisation, inductance calculation, coupling evaluation over 3D operational regions, and particle swarm optimisation under user-defined constraints. The toolchain is region-flexible: while demonstrated against SAE J2954 Z1, the evaluation grid can be redefined to target any arbitrary three-dimensional misalignment and air-gap envelope, extending the methodology's applicability to AGV chargers, docking robots, and other near-field WPT scenarios.
- **Systematic five-stage design pipeline:** A complete methodology comprising requirements definition, analytical pre-design and optimisation, high-fidelity FEA validation, ferrite and shielding integration, electrical worst-case assessment, and thermal simulation.
- **Experimental validation with error bounds:** Quantitative comparison between analytical predictions, FEA simulations, and experimental measurements, providing empirically grounded error margins for the proposed modelling approach.

1.6 Thesis Outline

The remainder of this thesis is organised into six chapters.

Chapter 2 presents the fundamental principles of inductive power transfer systems, covering magnetic induction, coil electrical properties, inductance calculations, resonant compensation, ferrite materials, and electromagnetic shielding. The chapter establishes the theoretical foundation for the toolchain developed in subsequent chapters.

Chapter 3 provides a critical review of the state of the art in IPT coupler design, covering coil geometries, ferrite structures, shielding techniques, misalignment handling, and the SAE J2954 standard. Gaps in the literature are identified to motivate the methodology introduced in this thesis.

Chapter 4 describes the MATLAB toolchain developed for rapid air-core pre-design of IPT couplers. The toolchain's ability to compute self-inductance, mutual inductance, and resistance is validated against FEA, confirming its reliability as a pre-design filter. The chapter concludes with a critical assessment of the toolchain's scope and future directions.

Chapter 5 presents the complete five-stage design pipeline, applying it to five coupler designs through optimisation, FEA validation, ferrite and shielding integration, electrical worst-case assessment, and thermal simulation. A critical assessment of the pipeline's limitations and future directions is provided.

Chapter 6 describes the practical implementation and experimental validation of three coupler designs, including the test jig fabrication, coil winding, and measurement procedures. The experimental results are compared with simulations, confirming the fundamental design ranking despite measurement challenges. A critical assessment of the experimental limitations and future directions is presented.

Chapter 7 concludes the thesis with a high-level synthesis of the main contributions. It reflects on what was achieved and where the work points next, and describes the dissemination of the MATLAB toolchain through an abstract accepted for poster presentation at the 10th E-Mobility Power System Integration Symposium, with the full conference paper currently in preparation. The final closing statement offers a forward-looking perspective on the continued evolution of the design pipeline.

Chapter 2

Fundamentals of Inductive Power Transfer Systems

2.1 System Architecture of an IPT Charger

Inductive power transfer has emerged as a leading wireless energy transfer technology for medium- to high-power applications, including Electric Vehicle (EV) charging, industrial automation, and mobile robotic platforms. By enabling contactless energy delivery through magnetic coupling between spatially separated coils operating at high frequency, IPT systems eliminate the need for physical electrical contacts. This characteristic significantly enhances robustness, operational safety, and mechanical flexibility, particularly in environments where frequent connection cycles, environmental exposure, or mechanical wear would otherwise degrade long-term reliability [1, 11, 17, 28, 63].

A generic IPT system can be decomposed into several functional stages, as illustrated in Fig. 2.1. Electrical energy from a primary source, typically the AC grid, is first conditioned by a power electronic front-end comprising rectification, filtering, and voltage regulation stages. A high-frequency inverter then generates an alternating current suitable for efficient magnetic coupling, which excites the transmitter-side compensation network and the Transmitter Coil (Tx), producing a time-varying magnetic field. Energy is transferred across an air gap to the Receiver Coil (Rx) via inductive coupling. On the receiver side, the induced alternating voltage is processed by a secondary compensation network, then rectified and filtered, and finally fed to a DC-DC converter to supply the load or charge an energy storage system.

The primary side (ground assembly) integrates the input filter and rectifier, the high-frequency inverter, the primary compensation network, and the Tx coil. The inverter typically operates at frequencies in the range of 79–90 kHz for EV charging applications, with 85 kHz being a common nominal value as recommended by the SAE J2954 standard [43]. The compensation network, which may be configured in series or parallel, cancels the inductive reactance of the coil, thereby reducing the reactive power drawn from the inverter and enabling high-efficiency power transfer. On the secondary side (vehicle assembly), the Rx coil is followed by a compensation network, a rectifier, a




Figure 2.1: Block diagram of a generic IPT system. Adapted from [9].

filter, and often a DC-DC converter that regulates the battery voltage and current according to the required charging profile.

Among these subsystems, the magnetic coupler – comprising the transmitter and receiver coils, ferrite structures, and electromagnetic shielding – plays a decisive role in determining overall system performance [22, 33]. While the power electronic stages can be actively controlled through modulation and feedback strategies, the magnetic coupling interface is largely defined by geometric design parameters, coil positioning, operating frequency, and magnetic material selection. Consequently, system performance is highly sensitive to spatial variations, air-gap changes, and application-specific environmental constraints.

Unlike conventional transformers that share a closed high-permeability core, IPT systems operate with relatively large air gaps (typically 100–250 mm for EV charging) and inherently weak magnetic coupling. This characteristic makes the design of the magnetic coupler particularly challenging, as performance indicators such as mutual inductance, coupling coefficient, and power transfer efficiency depend strongly on coil geometry, magnetic material distribution, and alignment conditions, especially under lateral misalignment and varying air-gap conditions [11, 24, 32, 56].

The SAE J2954 standard [43] defines operating frequency ranges, power classes, air-gap limits, and allowable misalignment tolerances for light-duty EV charging. It serves as a reference for the design of interoperable IPT systems. The following sections present the fundamental principles of magnetic induction, coil electrical properties, inductance calculations, resonant compensation, and the use of ferrite and shielding materials that underpin the design of such systems.

2.2 Magnetic Induction and Coupling

Inductive power transfer is fundamentally based on magnetic induction, whereby a time-varying current flowing through a conductor generates a corresponding time-varying magnetic field in the surrounding space. When a secondary conductor is placed within this magnetic field, an Electromotive Force (EMF) is induced according to Faraday’s law of electromagnetic induction. This mechanism enables wireless energy transfer between spatially separated coils without direct electrical contact [3, 56, 63].

In a basic IPT configuration, an alternating current in the transmitter coil generates a magnetic field that produces a magnetic flux, only a fraction of which links the receiver coil. Unlike

conventional transformers, IPT systems do not employ a shared high-permeability magnetic core, resulting in loosely coupled magnetic fields and substantial leakage flux. Consequently, only a portion of the magnetic flux generated by the Tx coil contributes effectively to power transfer. This effect is illustrated in Fig. 2.2, where the coils are shown along with the magnetic flux lines linking them.




Figure 2.2: Magnetic flux distribution in a loosely coupled IPT system. Adapted from [5].

The effectiveness of magnetic energy transfer is characterised by the mutual inductance M , which quantifies the amount of magnetic flux produced by the transmitter that links the receiver coil [11, 32, 34]. The self-inductances of the transmitter and receiver coils, denoted L_1 and L_2 , respectively, represent the magnetic energy stored locally in each winding. The voltage induced in the receiver coil can be expressed as (Eq. 2.1):

$$v_{Rx}(t) = M \frac{di_{Tx}(t)}{dt}, \quad (2.1)$$

highlighting the dependence of induced voltage on both coupling strength and the rate of change of the transmitter current [26, 47].

A normalised metric commonly used to characterise magnetic interaction is the coupling coefficient k , defined as:

$$k = \frac{M}{\sqrt{L_1 L_2}}, \quad (2.2)$$

with $0 \leq k \leq 1$. In practical IPT systems, k is typically low due to large air gaps and unavoidable misalignment, in contrast to conventional transformers where values approaching unity are achieved through the use of laminated magnetic cores [24, 40, 61].

Both mutual inductance and the coupling coefficient depend strongly on coil geometry, number of turns, conductor arrangement, relative positioning, air-gap distance, and the use of magnetic materials such as ferrites or flux guides. Spatial misalignment reduces magnetic flux linkage and degrades coupling, leading to reduced power transfer efficiency [10, 17, 37]. For example, for circular coils of equal radius r separated by a distance d (air gap), an approximate expression for

the coupling coefficient under perfect alignment is (Eq. 2.3):

$$k \approx \frac{r^3}{\sqrt{(r^2 + d^2)^3}}, \quad (2.3)$$

which shows that k decreases rapidly as the air gap increases. Under lateral misalignment Δx , the coupling coefficient further decreases since the overlapping area between the coils is reduced [38, 58].

Achieving high coupling coefficients in loosely coupled systems typically involves trade-offs in size, weight, and material cost [45, 62]. Increasing the number of turns or tightening conductor spacing may enhance inductance but also exacerbates the coil's equivalent resistance at high frequencies. Furthermore, larger coil diameters improve coupling at a given air gap but may be constrained by available space depending on the application. As a result, practical IPT systems for stationary charging, with air gaps of 100–150 mm, commonly exhibit coupling coefficients in the range of 0.2–0.4 [34, 38].

In summary, understanding magnetic induction, mutual inductance, and the coupling coefficient is essential for designing efficient IPT couplers. The following sections build on these concepts to discuss coil geometries, resonant compensation, and the role of ferrite and shielding materials.

2.3 Electrical Properties: Resistance, Skin and Proximity Effects

The efficiency of an IPT system depends not only on the strength of magnetic coupling but also on the electrical losses in the coil windings. When current flows through a conductor, power is dissipated as heat due to the conductor's resistance. Understanding the nature of this resistance – and how it changes with frequency – is essential for designing low-loss coils.

At Direct Current (DC) or very low frequencies, the resistance of a wire is given by the familiar formula:

$$R_{dc} = \frac{\rho l}{A}, \quad (2.4)$$

where ρ is the electrical resistivity of the conductor material (for annealed copper at 20 °C, $\rho = 1.72 \times 10^{-8} \Omega \cdot \text{m}$), l is the total length of the wire, and A is its cross-sectional area. For a given coil geometry, the wire length is determined by the number of turns and the average turn length. The DC resistance sets a lower bound on the winding losses. When an AC current flows through the wire, it will experience a higher resistance due to two distinct high-frequency phenomena: skin effect and proximity effect.

2.3.1 Skin Effect

Skin effect originates from Faraday's law of induction. When an alternating current flows through a conductor, the changing current produces a time-varying magnetic field both inside and around the wire. This magnetic field induces eddy currents within the conductor itself. According to Lenz's law, these eddy currents oppose the change in the original current. The result is a




Figure 2.3: Illustration of current distribution in a conductor: (a) DC, (b) skin effect at high frequency, (c) both skin and proximity effects combined. The colour scale ranges from low (blue) to high (red) current density. Adapted from [13].

non-uniform current density: the current is forced away from the centre of the wire and flows predominantly near the surface, as illustrated in Fig. 2.3b. The higher the frequency, the more pronounced this crowding becomes.

The skin depth δ is defined as the distance from the surface at which the current density falls to $1/e$ (approximately 37%) of its value at the surface. It is given by:

$$\delta = \sqrt{\frac{2\rho}{\omega\mu}}, \quad (2.5)$$

where $\omega = 2\pi f$ is the angular frequency, and μ is the magnetic permeability of the conductor (for non-magnetic copper, $\mu \approx \mu_0 = 4\pi \times 10^{-7} \text{ H/m}$). For copper at the typical IPT operating frequency of 85 kHz, the skin depth calculates to approximately 0.22 mm. This means that a solid copper wire with a diameter larger than about 0.44 mm (twice the skin depth) will have its inner portion underutilised, and the AC resistance will be significantly higher than the DC resistance. In practical IPT coils, which often use wire diameters of several millimetres, the skin effect alone would cause a substantial increase in losses if solid wire were used.

2.3.2 Proximity Effect

Even if skin effect is mitigated, another phenomenon called proximity effect can still affect the AC resistance. Proximity effect arises from the time-varying magnetic fields produced by adjacent turns of the same coil or by the other coil in the system. These external fields induce eddy currents in each conductor, further distorting the current distribution, as shown in Fig. 2.3c. In a tightly wound coil, the field from neighbouring turns is strong and non-uniform. Consequently, the proximity effect often dominates the total AC resistance, especially in multi-layer windings.

Unlike skin effect, which depends only on the wire's own current and operating frequency, proximity effect depends on the geometry of the winding, the spacing between turns, and the currents in nearby conductors. Quantitative prediction of proximity losses requires either finite element analysis or analytical models based on Dowell's method, which reduces a multi-turn

winding to an equivalent foil conductor. A detailed explanation of how Dowell's method is implemented in the present work is provided in Chapter 4.

2.3.3 Mitigation with Litz Wire

The combined consequences of skin and proximity effects can severely degrade IPT efficiency, since the AC resistance of a solid copper wire at 85 kHz may be several times its DC resistance. The standard mitigation is Litz wire (from German *Litzendraht*, meaning braided or stranded wire). Litz wire comprises many thin, individually insulated strands twisted together to form a bundle, which is then twisted along its length with a defined pitch (Figure 2.4). To suppress skin effect, each strand must have a diameter smaller than the skin depth at the operating frequency. The twisting pattern ensures that, on average, each strand occupies all radial positions within the bundle, equalising the magnetic flux linkage among strands. This equalisation cancels the eddy currents induced by the overall winding field, thereby drastically reducing proximity losses [27, 57].




Figure 2.4: Structure of a Litz wire: (a) 3-D model, (b) real Litz wire with visible bundle contours. Adapted from [13].

Properly designed Litz wire can achieve an AC resistance that is only slightly higher than the DC resistance of the bundle, even at frequencies of tens or hundreds of kilohertz. The design process involves selecting the strand diameter (based on the skin depth) and the number of strands (based on the required current-carrying capacity). For a given IPT application, these parameters are determined iteratively, as will be demonstrated in Chapter 5.

2.4 Electrical Properties: Self and Mutual Inductances

The performance of an inductive power transfer system is fundamentally determined by the self-inductances of the transmitter and receiver coils and by the mutual inductance that couples them. Accurate calculation or measurement of these quantities is therefore essential for designing the compensation networks and predicting the achievable efficiency.

2.4.1 Self-Inductance

Self-inductance L is a property of a single coil that quantifies its ability to store magnetic energy. For a coil with N turns, the self-inductance is defined as the ratio of the total magnetic flux linkage $\Lambda = N\phi$ (where ϕ is the flux through one turn) to the current i that produces it (Eq. 2.6):

$$L = \frac{\Lambda}{i}. \quad (2.6)$$

When the current changes, the induced voltage across the coil is $v = L di/dt$.

For a planar spiral coil (the most common geometry in IPT pads), closed-form analytical expressions provide good engineering estimates. One of the most widely used formulas is Wheeler's formula for a planar circular spiral (Eq. 2.7):

$$L(\mu\text{H}) = \frac{31.33 N^2 a^2}{8a + 11c}, \quad (2.7)$$

where $a = (D_o + D_i)/4$ is the mean radius (in metres), and $c = (D_o - D_i)/2$ is the radial build (the difference between outer and inner radii). As with the original Wheeler formula from which it is adapted, the leading numerical coefficient already absorbs the vacuum permeability and the unit conversion to microhenries, so μ_0 does not appear explicitly. The formula assumes that the turns are evenly spaced and that the coil is not too thick. For square or rectangular planar coils, analogous formulas exist, but they are more complex and often replaced by numerical integration or finite element analysis.

Alternatively, self-inductance can be computed analytically using the same Neumann integral employed for mutual inductance (Eq. 2.11), applied to a single closed contour. In this case, the integral must be evaluated with care to avoid singularities when the integration paths coincide. This approach is particularly useful when the coil geometry is arbitrary – not circular, not square, and not even necessarily a spiral. Although computationally more intensive than the approximation counterparts for known geometries, it allows for better parameterisation of the coil shape and does not incur any approximations. Hussain et al. [21] validated the discretised Neumann integral approach for planar spiral coils, showing good agreement with measurements.

2.4.2 Mutual Inductance

Mutual inductance M quantifies the magnetic coupling between two coils. It is defined as the ratio of the flux linkage in coil 2 due to a current in coil 1, or vice versa (Eq. 2.8):

$$M = \frac{\Lambda_{2 \leftarrow 1}}{i_1} = \frac{\Lambda_{1 \leftarrow 2}}{i_2}. \quad (2.8)$$

By reciprocity, the two definitions yield the same value. The induced voltage in coil 2 from a changing current in coil 1 is $v_2 = M di_1/dt$.

For simple geometries, closed-form approximations exist. Two coaxial circular loops of radii r_1 and r_2 , separated by a distance d , have a mutual inductance given by (Eq. 2.9):

$$M = \mu_0 \sqrt{r_1 r_2} \left[\left(\frac{2}{k} - k \right) K(k) - \frac{2}{k} E(k) \right], \quad (2.9)$$

where $K(k)$ and $E(k)$ are complete elliptic integrals of the first and second kind, and (Eq. 2.10)

$$k = \sqrt{\frac{4r_1 r_2}{(r_1 + r_2)^2 + d^2}}. \quad (2.10)$$

For planar spiral coils with many turns, the mutual inductance is often approximated by assuming that all turns are concentrated at the mean radius, or by using the filament method that sums the contributions of individual turns.

For arbitrary coil geometries – non-circular, non-square, or irregular shapes – mutual inductance can be computed exactly using the Neumann integral, derived from the vector potential form of Maxwell's equations [8]:

$$M = \frac{\mu_0}{4\pi} \oint_{C_1} \oint_{C_2} \frac{\mathbf{dl}_1 \cdot \mathbf{dl}_2}{|\mathbf{r}_1 - \mathbf{r}_2|}, \quad (2.11)$$

where C_1 and C_2 are the contours of the two coils, $\mathbf{dl}_1$ and $\mathbf{dl}_2$ are infinitesimal vector elements along the wires, and $\mathbf{r}_1$, $\mathbf{r}_2$ are the position vectors of these elements. The double contour integral accounts for the shape, orientation, and separation of the coils.

For numerical evaluation, each coil is discretised into a large number of small straight segments, as illustrated in Fig. 2.5. The double integral then becomes a double sum over all segment pairs:

$$M \approx \frac{\mu_0}{4\pi} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \frac{\Delta \mathbf{l}_{1,i} \cdot \Delta \mathbf{l}_{2,j}}{\|\mathbf{r}_{1,i} - \mathbf{r}_{2,j}\|}. \quad (2.12)$$

This discretised form is straightforward to implement in software and provides accurate results when the segment length is much smaller than the coil radius and the separation distance. For multi-turn coils, the total mutual inductance is obtained by summing over all turns, assuming that the current distribution is uniform.

The fundamental expressions above allow for the analytical estimation of self and mutual inductances for arbitrary coil geometries without approximations, although at a higher computational cost compared to closed-form formulas. This cost is justified when the goal is to study the parameterisation of coil shapes, as it enables systematic exploration of the design space and accurate coupling predictions under varying geometric conditions.

2.4.3 Experimental Determination of Inductance

In practical coil development, it is often necessary to measure self-inductance and mutual inductance rather than relying solely on calculations. Self-inductance can be measured directly using an Inductance, Capacitance, and Resistance (LCR) meter or an impedance analyser, applying




Figure 2.5: Discretisation of a planar spiral coil into small straight segments for numerical evaluation: (a) overall coil geometry, (b) discretised segments.

a test frequency close to the intended operating frequency. The measurement should be performed with the coil in its final assembled state, including ferrite if present, because ferrite significantly increases the inductance.

Mutual inductance can be determined by connecting the two coils in series and measuring the total inductance in two configurations: series-aiding and series-opposing, as shown in Fig. 2.6. If the two coils are connected so that their fluxes add, the total inductance is (Eq. 2.13):

$$L_{\text{aid}} = L_1 + L_2 + 2M. \quad (2.13)$$

If one coil is reversed, the fluxes oppose (Eq. 2.14):

$$L_{\text{opp}} = L_1 + L_2 - 2M. \quad (2.14)$$

Solving these two equations gives (Eq. 2.15):

$$M = \frac{L_{\text{aid}} - L_{\text{opp}}}{4}. \quad (2.15)$$

This method requires careful attention to the phasing of the coils (the polarity markings) but is straightforward and accurate. Pascoal [38] used this technique to validate his simulation results, achieving errors below 5% between measured and simulated mutual inductance.

2.5 Metrics for Coupling Performance

The preceding sections have established the fundamental electrical parameters of IPT coils: resistance (Section 2.3) and self- and mutual inductance (Section 2.4). From these parameters, two key metrics arise that collectively determine the coupling performance of an IPT system: the quality factor Q and the coupling coefficient k .




Figure 2.6: Measuring mutual inductance using the series-aiding (a) and series-opposing (b) method.

2.5.1 Quality Factor

The quality factor Q of a coil is a measure of its efficiency in storing magnetic energy relative to the energy dissipated as heat. It is defined as the ratio of the inductive reactance to the total resistance at the operating frequency (Eq. 2.16):

$$Q = \frac{\omega L}{R}, \quad (2.16)$$

where $\omega = 2\pi f$ is the angular frequency, L is the self-inductance, and R is the equivalent resistance of the coil. A high Q indicates that the coil stores energy effectively with low losses, which is essential for achieving high efficiency in resonant IPT systems.

The quality factor is frequency-dependent: it increases linearly with frequency (due to the ωL term) but eventually decreases when the resistance rises disproportionately due to skin and proximity effects. The choice of conductor therefore directly influences the achievable Q at the target operating frequency.

2.5.2 Coupling Coefficient

Recalling Eq. 2.2, the coupling coefficient k quantifies the magnetic coupling between the transmitter and receiver coils as the ratio of the mutual inductance M to the geometric mean of the self-inductances:

$$k = \frac{M}{\sqrt{L_1 L_2}}, \quad (2.17)$$

with $0 \leq k \leq 1$. In loosely coupled IPT systems, k values are low, often in the range of 0.1 to 0.4, due to the large air gap and unavoidable misalignment. The coupling coefficient depends on the coil geometry, relative positioning, air-gap distance, and the presence of ferrite or shielding materials. The dependencies and typical values of k were discussed in Section 2.2.

2.5.3 Implications for System Performance

Both Q and k are essential for predicting the achievable efficiency of an IPT system. For an SS-compensated system, the maximum efficiency is a function of the product $k^2 Q_1 Q_2$, as will be

shown in Section 2.6.3. This means that high efficiency requires both strong magnetic coupling (high k) and low coil losses (high Q).

The coupling coefficient is particularly sensitive to geometric variations. Under ideal conditions (air-core coils without ferrite or shielding), self-inductances L_1 and L_2 are essentially independent of the distance between coils and of lateral misalignment, because the flux that contributes to self-inductance is dominated by the local field around each coil. However, when ferrite or aluminium shielding is present, the mutual coupling between coils can affect the flux distribution in the ferrite, leading to a small variation in self-inductance with distance and misalignment. This effect is typically second-order but may become significant in highly coupled or tightly shielded systems. In contrast, mutual inductance M is always highly sensitive to both separation and misalignment, as it depends on the fraction of flux that links the other coil. Consequently, the coupling coefficient varies in the same manner as M . For a fixed coil geometry, M (and thus k) decreases monotonically as the air gap increases, and also decreases as lateral misalignment increases.

The quality factor, in turn, is strongly influenced by the choice of conductor. Using Litz wire is not merely a refinement but a necessity for IPT systems operating at 85 kHz and power levels of several kilowatts, as it reduces the AC resistance and thus maximises Q . The additional cost and manufacturing complexity of Litz wire are justified by the substantial reduction in winding losses and the corresponding improvement in system efficiency.

In summary, Q and k are the two primary metrics that define the coupling performance of an IPT system. Accurate computation of these metrics requires reliable models for resistance and inductance over the full operational envelope, including variations in air gap and misalignment. These concepts form the foundation for the design and optimisation methodology presented in the following chapters.

2.6 Resonant Operation and Compensation Networks

In a loosely coupled IPT system, the self-inductances of the transmitter and receiver coils are relatively large, while the mutual inductance is small. If the coils were connected directly to a voltage source, the inductive reactance ωL would cause a large phase shift between voltage and current, resulting in substantial reactive power circulation. This reactive current does not contribute to real power transfer but increases the RMS currents in the windings and the inverter, raising conduction losses and stressing the semiconductor devices.

To overcome this problem, compensation networks are added to the primary and secondary sides. These networks consist of capacitors connected in series or parallel with the coils, chosen such that the inductive reactance is cancelled at the operating frequency. When the system is correctly compensated, the input impedance becomes predominantly resistive, the reactive power circulation is minimised, and the power transfer capability for a given current is maximised.

2.6.1 Resonant Condition

For a single series LC circuit, resonance occurs when the inductive and capacitive reactances are equal in magnitude (Eq. 2.18):

$$\omega L = \frac{1}{\omega C}. \quad (2.18)$$

Solving for the angular frequency gives the resonant frequency (Eq. 2.19):

$$\omega_0 = \frac{1}{\sqrt{LC}}, \quad f_0 = \frac{1}{2\pi\sqrt{LC}}. \quad (2.19)$$

At this frequency, the impedances of the inductor and capacitor cancel, and the circuit behaves as a pure resistance. In an IPT system, both primary and secondary circuits are tuned to the same resonant frequency ω_0 , typically 85 kHz for EV charging applications as recommended by SAE J2954 [43].

2.6.2 Basic Compensation Topologies

Four fundamental compensation topologies exist, depending on whether the capacitor on each side is connected in series or parallel with the coil. Figure 2.7 illustrates these configurations.




Figure 2.7: Basic compensation topologies: (a) Series-Series (SS), (b) Series-Parallel (SP), (c) Parallel-Series (PS), (d) Parallel-Parallel (PP).

Series-Series (SS): Both capacitors are in series with their respective coils. The resonant frequency is load-independent, and the input impedance is purely resistive at resonance, simplifying control. For these reasons, SS is a widely used topology in static EV charging [28, 49, 62].

Series-Parallel (SP): The primary capacitor is in series with the Tx coil, while the secondary capacitor is in parallel with the Rx coil. This topology provides a current-source characteristic to the load, which can be beneficial for battery charging (constant current). However, the resonant frequency becomes load-dependent [62].

Parallel-Series (PS): The primary capacitor is in parallel with the Tx coil, and the secondary capacitor is in series with the Rx coil. This topology is less common because the primary side draws high circulating currents even at no load [62].

Parallel-Parallel (PP): Both capacitors are connected in parallel with their coils. Like the PS topology, the PP configuration suffers from high primary circulating currents and load-dependent tuning.

Beyond these four basic configurations, more advanced compensation networks have been developed to meet specific performance requirements. The double-sided Inductor-Capacitor-Capacitor (LCC) topology, which adds an additional inductor and capacitor in series with each coil, is widely used in high-power EV charging because it provides a constant current output independent of load and coupling, and it enables zero-voltage switching over a wide operating range [19, 54]. Other networks are also found in the literature, each offering different trade-offs in terms of component count, sensitivity to parameter variations, and misalignment tolerance [49]. The choice of compensation topology ultimately depends on the application requirements; for static EV charging with fixed coupling and moderate power levels, the simplicity and robustness of the SS topology remain attractive. The present work adopts SS compensation, which aligns with the power electronics design being developed in parallel within the same research group.

2.6.3 Deductions for the SS Topology

For the SS topology, the equivalent simplified circuit is shown in Fig. 2.8. The primary side consists of the inverter output voltage V_1 (modelled as a sinusoidal source at the resonant frequency ω_0), the primary compensation capacitor C_1 , the primary self-inductance L_1 , and a series resistance R_1 that represents the total resistance of the primary branch, mainly due to the coil resistance. The secondary side comprises the secondary self-inductance L_2 , its compensation capacitor C_2 , a series resistance R_2 (likewise dominated by the coil resistance), and an equivalent load resistance R_L . This resistance models the impedance seen from the secondary terminals as the combined effect of the rectifier, DC-DC converter, and battery. The mutual inductance M couples the two sides.




Figure 2.8: Equivalent circuit of an SS-compensated IPT system.

At resonance, the reactive impedances cancel:

$$\omega_0 L_1 = \frac{1}{\omega_0 C_1}, \quad \omega_0 L_2 = \frac{1}{\omega_0 C_2}. \quad (2.20)$$

Under this condition, the impedance reflected from the secondary to the primary is purely real, and the input impedance simplifies to (Eq. 2.21):

$$Z_{\text{ref}} = \frac{\omega_0^2 M^2}{R_L + R_2}, \quad Z_{\text{in}} = R_1 + Z_{\text{ref}}. \quad (2.21)$$

The efficiency η_{SS} can be derived as (Eq. 2.22):

$$\eta_{\text{SS}} = \frac{I_2^2 R_L}{I_1^2 R_1 + I_2^2 (R_2 + R_L)} = \frac{\omega_0^2 M^2 R_L}{(R_1 (R_L + R_2) + \omega_0^2 M^2) (R_L + R_2)}. \quad (2.22)$$

Setting the derivative with respect to R_L to zero yields the optimal load that maximises efficiency:

$$R_{L,\text{opt}} = R_2 \sqrt{1 + \frac{\omega_0^2 M^2}{R_1 R_2}} = R_2 \sqrt{1 + k^2 Q_1 Q_2}, \quad (2.23)$$

where $Q_1 = \omega_0 L_1 / R_1$ and $Q_2 = \omega_0 L_2 / R_2$ are the quality factors of the coils. Substituting this back gives the maximum efficiency for the SS topology [16]:

$$\eta_{\text{SS,max}} = \frac{k^2 Q_1 Q_2}{\left(1 + \sqrt{1 + k^2 Q_1 Q_2}\right)^2}. \quad (2.24)$$

This result is fundamental to IPT design: the maximum achievable efficiency depends only on the product $k^2 Q_1 Q_2$. Therefore, to achieve high efficiency, designers must maximise both the coupling coefficient (via coil geometry and alignment) and the quality factors (via low-resistance windings, typically using Litz wire).

The maximum efficiency derived in Eq. 2.24 is achieved when the equivalent load resistance seen by the secondary side is close to $R_{L,\text{opt}}$. In practice, this condition is not always met because the battery voltage and state of charge vary during the charging process. To address this, the DC-DC converter can be used as an impedance matching stage. By estimating the mutual inductance M from primary-side measurements, the converter adjusts its input impedance to emulate the optimal load resistance, thereby maintaining high efficiency across a range of operating conditions. This approach is often referred to as Maximum Efficiency Point Tracking (MEPT) and has been successfully implemented in IPT systems [50, 51, 60].

The above expressions for $R_{L,\text{opt}}$ and $\eta_{\text{SS,max}}$ are valid for the SS topology. For the other three basic topologies (SP, PS, PP) the same maximum efficiency expression holds under resonant conditions, but more complex networks such as LCC require different formulas.

2.6.4 Practical Aspects of Compensation

In practice, the compensation capacitors must be chosen to match the coil inductances at the nominal operating frequency. For the SS topology, the required capacitances are:

$$C_1 = \frac{1}{\omega_0^2 L_1}, \quad C_2 = \frac{1}{\omega_0^2 L_2}. \quad (2.25)$$

These capacitors must have low Equivalent Series Resistance (ESR) to avoid additional losses, and their voltage rating must exceed the peak voltages that appear across them, which can be significantly higher than the input voltage due to resonance.

The inverter output voltage is typically a square wave (due to switching operation), but the resonant tank filters out the higher harmonics, so the current is nearly sinusoidal. The fundamental component approximation is therefore valid for efficiency analysis.

Temperature variations, component tolerances, and changes in coil inductance due to ferrite saturation or proximity to metal can cause detuning – a shift away from the exact resonant frequency. Detuning increases reactive current and reduces efficiency. Therefore, practical IPT systems often include a control loop that adjusts the inverter frequency to maintain near-resonant operation [51, 60].

With the fundamentals of resonant compensation established, the next section discusses the role of ferrite materials in guiding magnetic flux and improving coupling.

2.7 Ferrite Materials for Flux Guidance

In an IPT system, the coils are separated by a relatively large air gap, and there is no closed high-permeability core to confine the magnetic field. Consequently, most of the magnetic flux produced by the transmitter coil does not reach the receiver coil, instead, it disperses into the surrounding space. This leakage flux not only reduces the mutual inductance and the coupling coefficient but can also induce unwanted eddy currents in nearby conductive structures.

To mitigate these problems, ferrite materials are placed behind or around the coils. Ferrites are ceramic compounds of iron oxide mixed with other metallic elements. They possess high magnetic permeability and very low electrical conductivity, making them ideal for guiding magnetic flux at high frequencies while suppressing eddy currents within the ferrite itself [55]. Unlike laminated steel cores used in conventional transformers (which require lamination to block eddy currents), ferrites are already electrically insulating, so they can be used as solid pieces without significant eddy-current losses.

2.7.1 Magnetic Properties of Ferrites

The key magnetic property of a ferrite is its relative permeability μ_r . A high-permeability ferrite can concentrate magnetic flux much more effectively than air. When placed behind a coil,

the ferrite provides a low-reluctance path for the magnetic field lines, pulling them through the ferrite and redirecting them towards the other coil.

However, ferrites are not ideal magnetic materials. They exhibit nonlinear behaviour: the permeability is constant only at low flux densities. As the magnetic flux density increases, the ferrite approaches saturation, typically at a saturation flux density B_{sat} of 0.3 to 0.5 T. Once saturated, the permeability drops sharply, and the ferrite ceases to guide flux effectively, behaving almost like air [55]. Therefore, ferrite cross-sections must be sized such that the peak flux density remains below B_{sat} under maximum operating current.

Ferrites also experience core losses, which consist of hysteresis losses (due to the energy required to realign magnetic domains) and eddy-current losses (induced circulating currents within the ferrite material). Because ferrites have high electrical resistivity, eddy-current losses are small, but hysteresis losses can still be significant at high frequencies and high flux densities [55]. This nonlinear behaviour is illustrated by the representative B-H curve in Fig. 2.9.




Figure 2.9: Nonlinear B-H characteristic of a typical ferrite material. Adapted from [14].

2.7.2 Common Ferrite Shapes for IPT Coils

Several ferrite geometries are used in IPT couplers. The choice depends on the coil shape, the available space, cost, and weight constraints. Figure 2.10 illustrates typical ferrite structures [33, 55].

Planar solid tiles. A planar tile is a single solid plate of ferrite that covers the entire area of the coil. This configuration provides the most effective flux guidance because the magnetic field lines have a continuous low-reluctance path. However, solid tiles are heavy, expensive, and brittle. Large tiles are difficult to manufacture and prone to cracking under mechanical stress [33].




Figure 2.10: Common ferrite structures: (a) planar solid tile, (b) segmented I-bars, (c) U-core, (d) E-core.

Segmented I-bars. To reduce weight and cost while retaining most of the flux-guiding benefit, segmented I-bars are often used. These are rectangular bars of ferrite placed radially under a circular coil or in parallel rows under a rectangular coil. The gaps between bars allow some flux to leak, but the overall coupling coefficient is only slightly reduced compared to a solid tile. The weight reduction can be significant, making I-bars the most common choice for practical EV charging pads [11, 38].

U and E cores. U-shaped and E-shaped ferrite cores provide a more closed magnetic circuit, similar to that of a conventional transformer. The coil is wound around the central leg of the core, and the outer legs return the flux. This configuration can achieve higher coupling coefficients for a given coil size, because the flux is confined almost entirely within the core. However, U and E cores are bulky, heavy, and sensitive to misalignment. They are more commonly used in stationary IPT systems where the coils are aligned precisely, rather than in EV charging where misalignment is expected [49, 61].

2.7.3 Effect of Ferrite on Inductance and Coupling

Adding ferrite increases both the self-inductances L_1 and L_2 and the mutual inductance M . The relative increase in M is typically larger than the increase in the self-inductances, because the ferrite channels the flux that links the two coils more effectively than the leakage flux that remains local. Consequently, the coupling coefficient $k = M/\sqrt{L_1 L_2}$ usually rises when ferrite is added, although the improvement is limited: once the ferrite already covers most of the coil area, adding more ferrite yields diminishing returns [10, 42].

The exact effect depends on the ferrite geometry and the air gap. For a given coil and air gap, the optimal ferrite size is often determined by simulation (finite element analysis) because the three-dimensional flux distribution is complex.

In summary, ferrite materials are essential for improving the coupling efficiency of IPT systems and for reducing stray magnetic fields. The choice of ferrite shape (solid tile, I-bars, U-core, or E-core) involves trade-offs among coupling improvement, weight, cost, and mechanical robustness. For most EV charging applications, segmented I-bars offer the best compromise. With the role of ferrites established, the next section discusses electromagnetic shielding, which complements ferrites by further containing stray fields to comply with safety standards.

2.8 Electromagnetic Shielding

Even with ferrite flux guides, a portion of the magnetic field generated by the transmitter coil leaks into the surrounding space. This stray field can induce eddy currents and voltages in nearby conductive objects, including the vehicle chassis, electronic control units, and metallic structures embedded in the ground. In addition, human exposure to time-varying magnetic fields must be kept within safety limits defined by the International Commission on Non-Ionizing Radiation Protection (ICNIRP) standards [30, 36]. For these reasons, IPT systems intended for commercial or public use require electromagnetic shielding to contain the stray field.

Shielding techniques for IPT systems can be classified into three categories: passive, active, and reactive shielding as illustrated in Fig. 2.11. Each method operates on a different principle and offers distinct trade-offs between effectiveness, complexity, and cost [25, 33, 47].




Figure 2.11: Classification of electromagnetic shielding techniques for IPT systems: (a) passive shielding, (b) active shielding, (c) reactive shielding. Adapted from [25].

2.8.1 Passive Shielding

Passive shielding is the simplest and most widely used method, particularly for medium-power EV charging systems. It consists of placing a conductive plate, typically made of aluminium or copper, on the side of each coil opposite the coupling gap. The plate is electrically isolated from the coil and is not connected to any external circuit.

When a time-varying magnetic field impinges on a conductive plate, eddy currents are induced within the plate according to Faraday's law. These eddy currents, in turn, generate a secondary

magnetic field that opposes the incident field. On the side of the plate away from the source, the incident and opposing fields partially cancel, thereby reducing the stray field. This behaviour is a direct consequence of Lenz's law: the induced currents act to oppose the change that created them.

The effectiveness of a passive shield depends on its electrical conductivity, thickness, and the frequency of the magnetic field. The skin depth δ in the shield material determines how deeply the field penetrates. For a shield to be effective, its thickness should be at least two to three skin depths at the operating frequency. For aluminium ($\sigma \approx 3.77 \times 10^7$ S/m) at 85 kHz, the skin depth is approximately (Eq. 2.26):

$$\delta_{Al} = \sqrt{\frac{2}{\omega \mu_0 \sigma}} \approx 0.28 \text{ mm.} \quad (2.26)$$

Thus, a 1–2 mm thick aluminium plate provides adequate shielding while being lightweight and cost-effective. Copper has higher conductivity but is more expensive and heavier, therefore, aluminium is the preferred material for passive shields in most IPT applications [47].

A practical consideration when adding a passive shield is its proximity to the coil. If the shield is placed too close (e.g., 1–2 mm), the eddy currents are strong and they also oppose the desired mutual flux, which can significantly reduce the coupling coefficient k and the self-inductances. If the shield is placed farther away (e.g., 20–30 mm), its effect on coupling is smaller, but the stray field containment is also less effective. In practice, a compromise distance of 10–20 mm is often chosen, and the reduction in k is accepted in exchange for meeting safety limits.

2.8.2 Active Shielding

Active shielding uses auxiliary coils, driven by external currents, to generate a magnetic field that cancels the stray field in specific regions. The auxiliary coils are typically placed around the periphery of the main coils or on the sides opposite the coupling gap. By controlling the amplitude and phase of the current in these auxiliary windings, the net magnetic field outside the coupling region can be reduced to near zero [25, 47].

The main advantage is that active shielding can achieve very low stray fields without significantly affecting the mutual coupling, as the auxiliary coils can be positioned to oppose only the stray component. It can also be dynamically adjusted to compensate for misalignment or load variations. However, active shielding adds significant complexity, cost, and weight, requiring additional windings, sensors, and control systems. It is typically reserved for high-power systems (above 100 kW) or applications with stringent electromagnetic compatibility requirements [18].

2.8.3 Reactive Shielding

Reactive shielding is a passive hybrid approach that uses a shorted conductive loop (or a resonant coil) placed in the stray field. The time-varying field induces a current in the loop, which generates a cancelling field without requiring an external power supply [25]. When the loop is tuned with a capacitor to resonate at the operating frequency, the induced current and cancellation effect are maximised [47].

Reactive shielding can be lighter and more compact than a solid aluminium plate, but it requires careful tuning: the resonant frequency depends on the loop inductance and capacitance, which must be stable under temperature variations. Moreover, the cancellation may not be optimal if the stray field changes due to misalignment. While promising, reactive shielding has not yet seen widespread commercial adoption in EV charging systems.

In summary, the ferrite and the aluminium shield work synergistically. The ferrite guides the flux towards the receiver, increasing coupling, while the aluminium shield contains the flux that would otherwise escape behind the coils. Together, they enable IPT systems to achieve high efficiency while complying with electromagnetic safety standards.

2.9 Finite Element Analysis and Thermal Simulation Tools

The analytical models presented earlier in this chapter are computationally efficient but rely on simplifying assumptions: thin filaments, air-core only, uniform current distribution, and the absence of nearby magnetic or conductive materials. For accurate prediction of electromagnetic and thermal behaviour in fully assembled couplers – incorporating ferrites and shielding – a more powerful numerical approach is required.

Finite element analysis solves Maxwell's equations in three dimensions over arbitrary geometries. The domain is discretised into small volume elements (tetrahedra or hexahedra), and the magnetic field, flux density, and induced currents are computed by solving the differential form of Maxwell's equations with appropriate boundary conditions. FEA can handle complex coil shapes, non-linear magnetic materials (ferrites), conductive materials with eddy currents (aluminium shields), and spatial variations such as misalignment, air-gap changes, and angular tilt. The main drawback is computational cost: a single 3D simulation can take minutes to hours, making exhaustive parametric sweeps impractical. The role that FEA plays alongside the analytical models introduced in this chapter, within the broader design methodology adopted in this work, is described in detail in later chapters.

2.9.1 Available FEA Software for IPT Design

Several commercial FEA packages are suitable for IPT coupler design [1, 15]. While COMSOL Multiphysics offers broad multi-physics coupling and JMAG is noted for efficient mesh generation, Ansys Maxwell is adopted in this work due to its widespread use in automotive WPT research and its seamless integration with Ansys Icepak for thermal co-simulation [38]. Maxwell is specialised for low-frequency electromagnetic problems and offers parametric sweeps, optimisation, and co-simulation with circuit simulators.

2.9.2 Ansys Maxwell for Electromagnetic Simulation

All electromagnetic simulations presented throughout this work use the Ansys Maxwell 2025 R2 frequency-domain (Eddy Current) solver. This solver is specifically suited for the 85 kHz

operating frequency, where skin and proximity effects are significant, but displacement currents remain negligible [38].

The dissertation does not aim to compare the relative performance of different FEA packages or solver formulations; rather, Ansys Maxwell is used as a validation and refinement tool within the broader design methodology. For the choice of mesh density, boundary conditions, and other simulation settings, this work follows the modelling conditions previously validated for IPT couplers by Pascoal [38], supplemented by the official Ansys Maxwell tutorials and documentation [5] where additional guidance was required.

In general terms, an adaptive tetrahedral mesh is employed, with manual refinement applied near coil windings and ferrite edges to resolve skin depth and flux gradients. The mesh refinement (Fig. 2.12) is concentrated in the air gap, around conductor surfaces, and at ferrite edges, ensuring accurate computation of mutual inductance and eddy current losses while maintaining acceptable computational cost. Loss data are extracted following standard Litz wire and ferrite core-loss modelling practices, the latter based on Steinmetz parameters derived from supplier datasheets [55].

The specific application of Maxwell within the multi-stage design pipeline developed in this thesis – including air-core validation, ferrite and shielding integration, and worst-case electrical assessment – is described in detail in Chapter 5.




Figure 2.12: Example comparison of (a) the initial tetrahedral mesh and (b) the adaptively refined mesh after several passes. Adapted from [5].

2.9.3 Ansys Icepak for Thermal Simulation

Copper losses in the windings and core losses in the ferrite generate heat. If the temperature exceeds the insulation class of the Litz wire or the rated temperature of the ferrite, the coupler may fail. Thermal simulation is therefore necessary to verify that the generated heat can be dissipated through conduction, convection, and radiation without exceeding material limits.

Ansys Icepak is a Computational Fluid Dynamics (CFD) tool integrated with the Ansys Electronics Desktop. It solves the coupled equations for fluid flow and heat transfer, importing loss data directly from Maxwell simulations and combining them with a thermal model of the assembled coupler. As with the electromagnetic simulations, the meshing strategy and boundary conditions adopted for the Icepak models presented in this work (natural convection coefficients, surface emissivity, ambient temperature) follow standard CFD practice as documented by the official Ansys Icepak tutorials [4], rather than being independently derived or benchmarked in this dissertation. The detailed thermal simulation workflow, including how loss data are transferred from Maxwell and how the resulting hotspot temperatures are interpreted, is presented alongside the corresponding stage of the design pipeline in Chapter 5.

In summary, Ansys Maxwell and Ansys Icepak are employed throughout this work as established, industry-standard tools for electromagnetic and thermal verification, rather than as objects of study in their own right. The dissertation's contribution lies not in the FEA or CFD methodology itself, but in how these tools are combined with the analytical toolchain to form an efficient, systematic design pipeline – a methodology introduced in Chapter 3 and developed in full in Chapters 4 and 5.

2.10 Summary

This chapter has presented the fundamental principles that govern inductive power transfer systems, with emphasis on the magnetic coupler. Starting from the basic mechanism of magnetic induction, the coupling coefficient k and mutual inductance M were introduced as the central figures of merit, followed by the frequency-dependent electrical behaviour of coils – skin and proximity effects, and their mitigation through Litz wire. Analytical and numerical methods for computing self- and mutual inductance were then described, including Wheeler's formula and the discretised Neumann integral, alongside the experimental series-aiding/series-opposing measurement technique. The chapter went on to cover resonant compensation networks, with particular attention to the series-series topology adopted throughout this work, and concluded with an overview of ferrite flux guides, electromagnetic shielding, and the FEA and thermal simulation tools used to validate fully assembled couplers.

Having established this theoretical foundation, the next chapter reviews the literature to examine how these principles have been applied and extended in published IPT coupler designs, while identifying the gaps that the methodology developed in this thesis aims to address.

Chapter 3

State-of-the-Art and Literature Review

3.1 Overview of IPT Systems

Inductive Power Transfer (IPT) has become the dominant wireless charging technology for stationary applications such as electric vehicles (EVs) and automated guided vehicles (AGVs), owing to its ability to transfer several kilowatts over large air gaps with efficiencies exceeding 90% [16, 28, 63]. A typical IPT system comprises a ground-side transmitter (Tx) coil, a vehicle-side receiver (Rx) coil, compensation networks, and power electronics. The magnetic coupler – the pair of coils together with ferrite flux guides and electromagnetic shielding – is the most critical subsystem because it determines the achievable coupling coefficient k and, consequently, the overall system efficiency [22, 33].

The field has been surveyed from multiple perspectives in the literature. Covic and Boys [16] provided a foundational treatment of IPT principles and coil topologies. Li and Mi [28] focused on EV applications, comparing compensation networks and safety standards. Zhang et al. [63] offered a broad review of near-field wireless power transfer (WPT), encompassing both inductive and capacitive methods. More recent contributions by Detka and Górecki [17] and Abuajwa et al. [1] have updated the state of the art, addressing advances in coil materials, misalignment mitigation, and efficiency benchmarks. Mohamed et al. [33] and Rahulkumar et al. [22] specifically examined inductive pad geometries – comparing circular, rectangular, Double-D (DD), and hybrid topologies – with respect to coupling performance, misalignment tolerance, and material cost. Shorter overviews are provided by Alphones and Jayathurathnage [3], Zamani et al. [61], and Liang [29].

Despite this progress, the literature reveals three persistent gaps: (i) most design optimisations consider only a single nominal alignment position, thereby neglecting off-nominal performance; (ii) ferrite and shielding are often selected heuristically rather than co-optimised with the coil geometry; and (iii) analytical modelling tools are either too simplistic for accurate loss prediction or too computationally expensive for systematic design-space exploration. The sections that follow critically examine these topics, identifying what the literature achieves, where it falls short, and how the present work addresses the identified limitations.

Before proceeding, it is important to acknowledge a fundamental challenge in interpreting

quantitative results from the literature: reported performance figures are strongly conditioned by the specific experimental or simulation context. Parameters such as power level, coil dimensions, air gap, misalignment range, ferrite configuration, operating frequency, and Litz wire construction vary widely across publications. Prototype power levels range from 50 W [40] to 11 kW [20]; air gaps span from 20 mm to 250 mm [23]; and operating frequencies range from 20 kHz [26] to 100 kHz [44]. Consequently, a “high efficiency” claim obtained under low-power, small-gap conditions cannot be directly extrapolated to an EV-scale application. The critical synthesis presented in this chapter therefore emphasises directional trends and qualitative advantages that are robust across a range of conditions, rather than absolute numerical comparisons, providing a more reliable foundation for the multi-stage design methodology introduced in this thesis.

3.2 Coil Geometries and Magnetic Coupler Topologies

A wide variety of coil geometries has been proposed for IPT systems, each presenting distinct trade-offs among coupling strength, misalignment tolerance, fabrication complexity, and material cost. This section critically reviews the principal topologies, highlights gaps in existing comparative studies, and positions the coil geometry choices made in this thesis. Fig. 3.1 shows the four topologies compared in this section.




Figure 3.1: Common IPT coil topologies (from left to right): Double-D Quadrature (DDQ), Double-D (DD), rectangular and circular. Adapted from [61].

3.2.1 Circular Coils

Circular coils are the most extensively studied IPT geometry owing to their rotational symmetry, analytical tractability, and ease of manufacture. An example circular coil and the field distribution is presented in Figure 3.2. Budhia et al. [11] demonstrated that circular pads combined with ferrite flux guides can achieve peak efficiencies above 95% for a 2 kW, 20 kHz system at a 210 mm air gap. However, their coupling coefficient degrades rapidly under lateral misalignment: a 50 mm displacement can reduce k by more than 40% [6, 23]. For stationary EV charging, where parking tolerances can be substantial, circular coils are therefore rarely adequate without active alignment systems or supplementary misalignment-tolerant compensation strategies.

On the positive side, circular coil design is well understood analytically, and several optimisation studies [6, 58] provide clear design rules governing turn count, wire gauge, and outer diameter. A limitation of most of these rules, however, is that they are derived at a fixed operating frequency and a specific ferrite layout, which limits their generalisability. Nataraj et al. [34] compared

circular, square, and rectangular coils and reported that circular coils yield the highest mutual inductance and coupling factor for a given footprint; their study employed small-scale, low-power coils without ferrite, however, so the conclusions may not scale to EV charging pads. Kosesoy et al. [26] implemented a 1 kW WPT system using circular pads at 20 kHz and achieved 89% coil-to-coil efficiency. This figure would differ at the SAE J2954 nominal frequency of 85 kHz due to skin and proximity effects. Jaafari et al. [23] demonstrated that enclosing circular coils in ferrite boxes can increase k by approximately 50% (from 0.38 to 0.58 at a 140 mm air gap) and improve misalignment tolerance, albeit at higher cost and mass.




Figure 3.2: Example of a circular coil geometry (a) and the corresponding magnetic field vector distribution in the YZ plane (b).

A notable limitation of most circular-coil studies is that optimisation targets a single nominal air gap and centred alignment, without evaluating mean coupling performance over the full misalignment envelope. The present work addresses this by applying spatial averaging over the entire operational region defined. This reveals that the high peak efficiency characteristic of circular coils often comes at the expense of substantially degraded off-nominal performance.

3.2.2 Rectangular and Square Coils

Rectangular and square coils offer improved misalignment tolerance along their longer axis, because the effective overlap area between transmitter and receiver diminishes more slowly under displacement in that direction. Ongayo and Hanif [37] used finite element analysis (FEA) to demonstrate that rectangular coils maintain higher mutual inductance than circular coils for misalignments up to 75 mm. Their study was performed at 20 kHz and without ferrite, so the quantitative results would differ at 85 kHz owing to altered eddy-current and proximity-effect losses. A known drawback of rectangular geometry is non-uniform flux density, with the coil corners exhibiting elevated leakage and potential thermal hotspots. Figure 3.3 illustrates the magnetic field

distribution generated by a rectangular coil, revealing field hotspots at the inner corners that are expected to manifest as thermal hotspots.




Figure 3.3: Example of a rectangular coil geometry (a) and the corresponding magnetic field vector distribution in the YZ plane (b).

Systematic optimisation of the coil aspect ratio for a given misalignment envelope has received comparatively little attention in the literature. Chatterjee et al. [12] presented a WPT design optimisation approach but focused primarily on circular topologies using analytical approximations that neglect ferrite saturation and proximity losses. Herceg et al. [20] designed double-layer rectangular coils for an 11 kW, 85 kHz WPT system and showed that multi-layer windings improve power density at the expense of increased manufacturing complexity. Their study is directly relevant to the present work given the shared frequency and power class. Solouma et al. [46] introduced a quasi-elliptical parameterisation that smoothly interpolates between circular and rectangular geometries, optimising the coupling coefficient via the Neumann integral – an approach conceptually aligned with the analytical toolchain developed in this thesis. Their optimisation, however, was performed at low power and without ferrite.

The literature correctly identifies the aspect-ratio trade-off: longer coils tolerate displacement along the elongated axis but become more sensitive to transverse misalignment. What is notably absent are quantitative guidelines for selecting the optimal aspect ratio based on the anticipated distribution of misalignment (e.g., preferential longitudinal versus lateral parking error). The present work addresses this through a parametric study of rectangular coils with varying aspect ratios, optimised for any coupling metric across the full operational region rather than merely the peak value.

3.2.3 Double-D (DD) and Related Topologies

Double-D (DD) coils consist of two rectangular windings placed side by side and connected in series with opposing polarity, such that their magnetic flux contributions reinforce one another in

the intermediate region. The resulting flux path traverses the air gap in a predominantly horizontal plane, conferring superior misalignment tolerance along the axis perpendicular to the winding split. This design was pioneered by Budhia et al. [11] and subsequently adopted as a reference pad geometry in SAE J2954. Fig. 3.4 shows a representative DD coil geometry together with its simulated field pattern.

The magnetic field distribution of a DD coil, shown in Fig. 3.4b, illustrates the characteristic horizontal flux path between the two D-shaped windings. While this configuration improves lateral misalignment tolerance, it also introduces asymmetric flux patterns that can either comply with or violate human electromagnetic exposure limits depending on orientation. Additionally, the field concentration at the corners and outer edges indicates regions of high localised flux density that are expected to correspond to elevated ohmic losses and, consequently, thermal hotspots in the physical implementation.




Figure 3.4: Example of a DD coil geometry (a) and the corresponding magnetic field vector distribution in the YZ plane (b).

Mohamed et al. [32] demonstrated that DD coils can retain 90% of their nominal coupling coefficient up to 100 mm of lateral misalignment when paired with appropriately configured ferrite. Their optimisation employed 2-D FEA and a specific turn-spacing trade-off that may not generalise to different pad sizes or air gaps. Zamani et al. [61] identified DD pads as among the most promising candidates for high misalignment tolerance, though their review does not offer cross-topology comparisons normalised for equal ferrite volume or copper area. Prosen et al. [41] extended the DD concept to vertically stacked double-DD configurations enabling independent power delivery to two receivers; their prototype operated at low power (approximately 50 W) and reduced pad dimensions, leaving the coupling behaviour at EV scale uncertain.

However, DD coils have significant drawbacks beyond the field distribution issues already noted. They require precise ferrite placement – typically two bar-shaped ferrites beneath each rectangular winding – and exhibit pronounced sensitivity to angular (yaw) misalignment. Furthermore, most

comparative studies evaluate DD coils only against circular pads rather than against properly optimised rectangular designs, making it difficult to isolate the performance contribution of the DD winding configuration from differences in coil area or ferrite volume. Rahulkumar et al. [22] concluded that DD pads offer the best misalignment performance, but at the highest material cost. A rigorous apples-to-apples comparison accounting for equal copper cross-sectional area, equal ferrite volume, and equal air gap is conspicuously absent from the literature.

3.2.4 Hybrid and Multi-Coil Designs

Hybrid topologies – including Double-D Quadrature (DDQ), bipolar pads, and square-hexagonal geometries – have been proposed to further extend misalignment tolerance or enable multi-directional coupling [24, 40]. Kalwar et al. [24] demonstrated that DDQ pads reduce null coupling zones relative to standard DD, though their experimental validation was conducted at 33 kHz and 700 W, far below the 85 kHz, multi-kW regime of interest for EV charging. Pashaei et al. [40] introduced a square-hexagonal hybrid coil, reporting a 17% increase in k over a conventional square coil at 50 mm air gap. However, their 50 W prototype lacked ferrite shielding, so the loss behaviour and leakage field distribution at full EV scale would differ substantially. Awuah et al. [7] presented a rectangular coil design with spatially varying trace pitch, width, and thickness, achieving a 19.24% improvement in quality factor at 85 kHz without increasing the coil footprint – a result directly relevant to the AC loss modelling developed in this thesis. Their optimisation targeted a PCB-style coil without ferrite, however, so the gain within a fully shielded EV pad remains to be verified.

While conceptually innovative, hybrid designs introduce substantial fabrication complexity: multiple winding layers, non-planar ferrite structures, and tight manufacturing tolerances. More importantly, their performance advantage is most often demonstrated at a single air gap or under idealised conditions (low power, no ferrite, no misalignment). The literature lacks systematic studies that quantify the marginal benefit of hybrid topologies over well-optimised DD or rectangular designs when manufacturing imperfections and full-envelope misalignment are accounted for.

3.3 Ferrite Structures and Electromagnetic Shielding

Ferrite materials are nearly ubiquitous in modern IPT couplers due to their ability to concentrate magnetic flux, increase mutual inductance, and reduce stray fields. Likewise, aluminium shielding is often added to contain leakage flux and meet electromagnetic safety standards. This section critically reviews the state of the art in ferrite and shielding design, highlighting heuristic practices, existing gaps, and the approach adopted in this thesis.

3.3.1 Ferrite Arrangements: From Heuristics to Optimisation

Common ferrite geometries include planar tiles, segmented bars, U-cores, E-cores, and I-cores, as reviewed by Wang et al. [55] and Mohamed et al. [33]. Wang et al. also compare conventional

ferrites with emerging materials such as nanocrystalline alloys and metamaterials. Planar tiles are the simplest to manufacture and assemble but provide limited flux guidance. Segmented bars – for example, radially arranged strips beneath a circular coil – offer a better trade-off between flux concentration and mass [11]. U- and E-cores establish a more closed magnetic circuit, approaching transformer-like behaviour, but are heavier, more expensive, and more sensitive to misalignment [49]. Shevchenko et al. [44] observed that ferrite shape significantly affects coupling, though their study was conducted at low power (50 W) and relatively high frequency (100 kHz), so the quantitative conclusions may not scale directly to EV systems.

Several FEA-based comparative studies have examined different ferrite configurations for given coil geometries. Bensetti et al. [10] optimised ferrite structures for dynamic WPT, finding that segmented bars reduce mass while preserving coupling, that flux concentration saturates beyond a certain ferrite surface area, and that a central aperture of up to 150 mm does not appreciably degrade coupling. Their work focused on dynamic charging with rectangular coils, which differs from the static charging scenario considered here. Henriques et al. [31] proposed concave ferrite cores to improve misalignment tolerance, but their achieved efficiency of 25.4% is far below EV requirements, rendering the concept interesting in principle but not yet practically viable. Rhee et al. [42] investigated the role of ferrite permeability in low-power WPT, demonstrating that efficiency increases with permeability but saturates beyond approximately 600; using materials with permeability of 3200 or higher adds mass and cost with negligible efficiency gain. Although conducted at power levels (160 mW) orders of magnitude below EV charging, this saturation principle is nonetheless instructive for material selection.

An important and frequently overlooked consideration is the influence of ferrite on coil self-inductance. Pascoal et al. [39] showed that adding ferrite plates increases self-inductance, thereby shifting the resonant frequency away from the design point; the compensation network must consequently be retuned. This effect is often ignored in studies that treat ferrite as a simple performance booster.

Despite these contributions, ferrite design is commonly performed after the coil geometry has been fixed rather than systematically explored across candidate layouts. The choice of geometry (tile, bar, or E-core) is frequently made heuristically based on precedents from SAE J2954 without quantitative justification. Moreover, prior work within the same research group [38, 50] relied on manual evaluation of only a small number of ferrite configurations. The present work follows a similar approach – a set of candidate ferrite layouts are defined for the coil geometries produced in Stages 1 and 2 of the design pipeline (Chapter 5) – but extends the evaluation methodology. Each layout is combined with a fixed aluminium backplate for passive shielding, and the candidate configurations are assessed using a cost-performance metric based on the maximum expected coil-to-coil efficiency computed across a misalignment region rather than solely at the aligned position.

3.3.2 Shielding Techniques and Their Trade-Offs

Aluminium plates are the most widely used passive shields in IPT systems. Positioned behind the Tx and Rx coils, they attenuate stray magnetic fields in the downward (ground) and upward (vehicle) directions [25, 47]. Active shielding using coils driven in phase opposition has been demonstrated but adds system complexity and cost. Reactive shielding using shorted loops offers a passive alternative, though it is rarely found in commercial systems. Foote and Onar [18] concluded that passive aluminium shielding remains the most practical solution for medium-power EV chargers; Covic and Boys [16] acknowledge active shielding as viable at high power but do not provide detailed design rules or cost-benefit analyses.

Tan et al. [47] compared ferrite, aluminium, and resonant reactive current (RRC) shields for a 5 kW system, finding that RRC shielding reduces the magnetic field strength by 81–95% while preserving coupling, though it requires careful tuning of the resonant loop. Kim et al. [25] introduced the reactive shield concept and characterised its near-field cancellation mechanism, largely through simulation without experimental verification. It is well established in the literature that shield proximity affects coil inductance and quality factor: a shield placed within 1–2 mm of the coil significantly opposes mutual flux, reducing k .

Nevertheless, many publications treat shielding as an afterthought, adding a standard 2 mm aluminium plate without evaluating its impact on efficiency or optimising its thickness and stand-off distance. Quantitative trade-off studies relating shield thickness and position to both leakage field reduction and efficiency penalty are scarce. The present work includes a parametric analysis of aluminium shielding (Chapter 5) to assess the trade-off between coupling coefficient degradation and shield thickness and stand-off distance, considering the practical constraints of mechanical integration and material cost.

3.3.3 Critical Synthesis: Why Co-Design Matters

A recurring theme in the reviewed literature is the sequential, component-wise design of coils, ferrites, and shields. Each element is optimised in isolation, leading to suboptimal system-level performance. For example, a ferrite arrangement that maximises k at centred alignment may perform poorly under misalignment, and an aggressively thick shield may introduce excessive eddy-current losses. Vaka and Keshri [49] emphasised the need for integrated design approaches but did not propose a concrete methodology. Prior work within the same research group [38, 50] carried out manual parametric sweeps of ferrite configurations; Viera's later contribution [54] extends the optimisation to impedance-plane interoperability but still treats ferrite and shielding as separate design steps.

The present work addresses this gap through a staged pipeline in which coil geometry is first resolved analytically and then validated by FEA before ferrite and shielding are introduced. The analytical toolchain (Chapter 4) does not model ferrite, but it rapidly identifies coil geometries with robust spatial coupling performance. Only these screened candidates proceed to the next stage, where a defined set of ferrite tile layouts – each combined with a fixed aluminium backplate

– is similarly analysed in FEA software. Layout selection is based on the maximum expected coil-to-coil efficiency evaluated over the defined misalignment region, providing a performance assessment that is representative of real operating conditions rather than the single aligned position used in most prior work. This staged approach reduces the total number of FEA evaluations required while ensuring that the chosen ferrite configuration is matched to the spatial coupling characteristics of the coil geometry it accompanies.

3.4 Misalignment and Coupling Variability

In practical IPT installations, perfect alignment between Tx and Rx coils cannot be guaranteed. Lateral displacement, vertical air-gap variation, and angular tilt all degrade mutual inductance and the coupling coefficient. This section reviews how the literature quantifies coupling variability and identifies critical gaps in the prevailing methodology. Nguyen et al. [35] surveyed electromagnetic-field-based WPT for unmanned aerial vehicles, confirming that misalignment is a pervasive challenge across all near-field WPT applications.

3.4.1 Coupling Degradation Under Misalignment

Numerous experimental and numerical studies have characterised how k degrades with misalignment, though the specific displacement ranges and air gaps considered vary considerably. Yang et al. [58] reported that for square coils, a 75 mm lateral displacement reduces k by 50–60%, whereas increasing the air gap from 100 mm to 200 mm reduces k by approximately 30%. Jaafari et al. [23] optimised circular coils with ferrite boxes and observed that lateral misalignment beyond 100 mm substantially degrades k even in the presence of ferrite. Liang [29] identified misalignment tolerance as a key differentiator among commercial systems. Aslan et al. [6] presented parametric sweeps of circular coils and efficiency-versus- k curves for the series-series (SS) compensation topology. Kosesoy et al. [26] and Shi et al. [45] also reported misalignment sweeps, but at 20 kHz, so their quantitative results are not directly transferable to 85 kHz systems. The qualitative effect of lateral and vertical misalignment on coupling performance is illustrated in Fig. 3.5.

Many studies include parametric sweeps of misalignment, generating 2-D contour plots of k or efficiency that provide valuable sensitivity insight. However, these sweeps are typically limited to a handful of discrete displacement values (e.g., 0, 50, 100 mm) rather than a dense grid covering the full operational region, and the ranges studied vary widely across publications. Crucially, few studies report statistical summaries – mean, minimum, or standard deviation of k – over a predefined operational envelope. The absence of such metrics makes fair cross-topology comparison difficult: a design with a higher peak k may exhibit a substantially lower spatial average, which is more representative of real-world performance.

The present work introduces a spatial averaging methodology that computes the mean, minimum, and coefficient of variation of the mutual inductance M and consequently the coupling coefficient k over a dense grid of misalignment positions covering the prescribed operational region. This enables a robust, statistically grounded comparison of coil topologies and guides the




Figure 3.5: Comparison of coil performance under misalignment conditions: (a) lateral misalignment and (b) vertical misalignment. Adapted from [58].

optimisation toward designs that perform consistently well across the expected range of parking positions.

3.4.2 Implications for Efficiency and Control

Coupling variability has direct consequences for system efficiency. From the optimal load expression for the SS topology (see Eq. 2.23 in Chapter 2), the maximum achievable efficiency is

$$\eta_{SS,max} = f(k^2 Q_1 Q_2),$$

where Q_1 and Q_2 are the quality factors of the transmitter and receiver coils, respectively. A reduction in k from 0.3 to 0.15 decreases the maximum achievable efficiency from approximately 95% to 85%, even when the load impedance is re-optimised. Under a fixed load – as occurs when the battery converter directly terminates the secondary – the efficiency penalty can be far more severe. Yin et al. [60] proposed a primary-side method for estimating mutual inductance and load resistance from input measurements, enabling adaptive control without secondary-side communication. The development of such control algorithms is outside the scope of the present thesis; the coupler design methodology developed here is nonetheless intended to supply the control engineer with reliable statistics of the expected k range over the operational envelope, so that appropriate tracking strategies can be designed accordingly.

A recurring concern in the literature is that efficiency is most often reported at a single, perfectly aligned condition, with claims frequently exceeding 95%. These figures are not representative of real-world operation in the presence of misalignment. Without systematic evaluation across the full operational region, such claims may present an optimistic and potentially misleading picture of practical system performance.

3.5 SAE J2954 Standard: Critical Analysis

SAE J2954 [43] is the primary international standard for wireless power transfer in light-duty EVs. It defines power classes, Z-classes for ground clearance, operating frequency range (79–90 kHz, nominal 85 kHz), and misalignment tolerances. It also specifies interoperability classes, electromagnetic safety limits, and minimum system efficiency requirements. This section critically examines what the standard provides, what it omits, and how the present work uses it as a boundary condition rather than as a prescriptive design guide.

3.5.1 What the Standard Provides

SAE J2954 is an indispensable reference for ensuring interoperability between vehicles and charging infrastructure from different manufacturers. It establishes a common terminology for power levels, air gaps, and test procedures, and it incorporates electromagnetic safety limits consistent with ICNIRP guidelines [30, 36]. The standard also mandates foreign object detection and living object protection.

Two interoperability classes are defined. Class I applies to public charging infrastructure, which must function across a wide range of vehicle ground clearances and misalignment conditions. Class II applies to dedicated systems in which the ground assembly (GA) and vehicle assembly (VA) are designed jointly for a specific application.

The standard's power classes (WPT1–WPT3) and Z-classes (Z1–Z3), which define the vertical separation between the VA coil and the GA coil, are summarised in Table 3.1 and Table 3.2, respectively.

Table 3.1: SAE J2954 wireless power transfer (WPT) power classes.

Class	Apparent power
WPT1	0–3.7 kVA
WPT2	0–7.7 kVA
WPT3	0–11.1 kVA

Table 3.2: SAE J2954 Z-classes for ground clearance (VA to GA).

Class	Distance range
Z1	100–150 mm
Z2	140–210 mm
Z3	170–250 mm

The misalignment types and their specified tolerances are summarised in Table 3.3. The coordinate axis definition (X: vehicle longitudinal direction; Y: vehicle lateral direction; Z: vertical) is illustrated in Fig. 3.6.

Table 3.3: SAE J2954 misalignment tolerances.

Misalignment type	Tolerance
Longitudinal (X)	± 75 mm
Lateral (Y)	± 100 mm
Vertical (Z)	Z1–Z3 ranges (Table 3.2)
Roll	$\pm 2^\circ$
Pitch	$\pm 2^\circ$
Yaw	$\pm 3^\circ$




Figure 3.6: Coordinate axes for SAE J2954 misalignment definitions. Adapted from [43] and [38].

The standard also specifies minimum system efficiency requirements. For Class I interoperability, the efficiency requirements, which depend on both WPT class and alignment condition, are shown in Table 3.4. For Class II systems the efficiency must be $\geq 85\%$ at the centred position and $\geq 80\%$ over the full range of misalignment and air-gap variations.

Table 3.4: SAE J2954 Class I minimum system efficiency requirements.

WPT Class	Centred position	Alignment area
WPT1	80%	75%
WPT2	82%	77%
WPT3	85%	80%

The standard further addresses inter-power-class interoperability, requiring that a GA rated for a higher power class (e.g., WPT3) must be capable of charging a VA rated for a lower class (e.g., WPT1). This imposes additional constraints on the magnetic coupler design to ensure adequate coupling across non-matched coil geometries. Electromagnetic field safety for human exposure, in accordance with ICNIRP guidelines [30, 36], is also an integral requirement, particularly for Class I interoperability.

3.5.2 Critical Analysis: What the Standard Does Not Specify

While SAE J2954 is comprehensive as a performance specification, it is explicitly not a design guide. Several significant gaps remain.

3.5.2.1 Lack of Justification for Coil Geometries

The most striking deficiency is the standard's treatment of coil geometry. Normative Appendix A (Test Station VA Specifications) and Appendix B (Test Station Universal GA) specify mechanical and electrical characteristics for all power and Z-class combinations for Class I interoperability. However, these specifications are presented in a disjointed, non-uniform manner that suggests compilation from contributions of different working groups rather than a coherent, unified exposition.

For Class I systems, the Universal GA is specified as a rectangular coil with rounded corners, while the Test Station VAs are square coils with rounded corners – neither strictly rectangular nor circular. Detailed mechanical dimensions (outer length, width, corner radius, turn count, Litz wire construction) and electrical specifications (inductance ranges, compensation topology) are provided for each combination; the rationale behind these specific geometries, however, is never explained. It remains unclear whether the choices were driven by efficiency, cost, misalignment tolerance, interoperability, or some weighted combination.

The disparity in documentation depth is illustrated in Fig. 3.7: the Universal GA specification is highly detailed, including precise corner radii, turn counts, Litz wire specifications, and ferrite placement. The Test Station VA drawing, by contrast, shows only the coil outline with fewer construction details. This inconsistency suggests that different sections of the standard were developed by different working groups with varying documentation standards.

3.5.2.2 Additional Gaps

Beyond the coil geometry issues, several other gaps remain:

- **Absence of optimisation metrics:** The standard does not define quantitative design objectives for the magnetic coupler (e.g., minimum acceptable spatial mean of k over the misalignment envelope, maximum permissible ferrite mass). As a result, researchers and manufacturers adopt arbitrary targets, complicating cross-study comparison.
- **Limited guidance for scaling:** Extensions to higher power classes or non-automotive applications are left entirely to the designer. The reference geometries for WPT1–WPT3 may not scale linearly in coupling or efficiency.
- **No ferrite optimisation procedure:** The standard illustrates example ferrite layouts but provides no guidance on how to adapt them for different coil dimensions or air gaps.

As noted by Pascoal [38], while the standard specifies the misalignment envelope and efficiency targets, it does not prescribe how to achieve them. The disjointed presentation of coil geometries, combined with the lack of coupling estimates, leaves a significant gap between the specification and the actual design of the magnetic coupler.

3.5.2.3 A Path Forward: Replication and Validation

One particularly valuable research direction is the exact replication of the standard's reference designs – following the specified mechanical dimensions, turn counts, Litz wire construction, and ferrite layouts – and their experimental characterisation for system efficiency. Such empirical validation would be of considerable benefit to the research community and to the standard itself. This work is currently being pursued by Viera's team [54], and the forthcoming results are expected to clarify the actual performance of the standard's reference geometries and to identify areas where the standard could be improved or made more internally consistent.





Figure 3.7: Example coil geometries from SAE J2954 appendixes: (a) universal GA (for all power and Z-classes) and (b) test station VA for WPT1/Z1 class. Adapted from [43].

3.5.3 How the Present Work Uses the Standard

The present work uses SAE J2954 in the following ways, without slavishly copying its reference geometries:

1. **Operational envelope:** The misalignment range (± 75 mm in X, ± 100 mm in Y) and the air-gap range of the Z1 class (100–150 mm) define the spatial domain over which coupling statistics are evaluated (Chapter 4).
2. **Power class:** The design targets the WPT1 class (3.7 kVA) as a representative medium-power level for residential and public charging scenarios.
3. **Operating frequency:** The nominal frequency of 85 kHz is adopted to ensure compatibility with standard-compliant components and verification equipment.
4. **Design freedom:** Coil, ferrite, and shield geometries are optimised from first principles within the standard's dimensional and performance constraints, rather than constrained to the standard's example pads. This approach permits the exploration of configurations that may outperform the reference designs while remaining fully compliant.

3.6 Design Optimisation Approaches

A range of methods has been applied to the design and optimisation of IPT magnetic couplers. This section critically reviews the most prevalent approaches, highlighting their respective strengths and limitations. Xuan Bac et al. [9] demonstrated a SiC-based matrix converter topology for IPT, focusing on power electronics rather than coupler optimisation. Zavrel and Kindl [62] compared compensation topologies (SS, SP, PS, PP) for EV charging, concluding that the series-series (SS) configuration is the most suitable for high power and misalignment tolerance.

3.6.1 Finite Element Analysis: Accurate but Computationally Intensive

FEA (implemented, for example, in Ansys Maxwell, COMSOL Multiphysics, or JMAG) numerically solves Maxwell's equations for arbitrary geometries, including nonlinear ferrite behaviour and complex 3-D coil shapes. It is the gold standard for accuracy and is widely used throughout the literature [6, 10, 11, 20, 23, 32, 40, 58].

Strengths: FEA handles ferrite saturation, eddy-current losses in shields, and realistic 3-D geometries with high fidelity. It produces flux density maps, volumetric loss distributions, and accurate coupling coefficients.

Limitations: FEA is computationally expensive. A single 3D simulation can take minutes to hours, making it impractical for exhaustive parametric sweeps or evolutionary optimisation with hundreds of evaluations. Many studies therefore limit the number of design variables or evaluate only a handful of candidate designs. Relying solely on FEA encourages a “design then simulate” approach rather than true optimisation. Researchers often simulate a few intuitive designs

and declare the best one optimal, but without thorough exploration they cannot guarantee global optimality.

3.6.2 Analytical Formulas: Fast but Limited in Scope

Closed-form expressions for mutual inductance and self-inductance exist for simple coaxial geometries (circular and rectangular loops), derived from the Neumann integral or empirical approximations [8, 21, 45]. These expressions evaluate almost instantaneously, enabling rapid parametric sweeps over large design spaces. Solouma et al. [46], for example, used the Neumann integral to optimise quasi-elliptical coil shapes via brute-force search; Pascoal [38] implemented a discretised Neumann integral in MATLAB to compute inductances for circular coils.

Strengths: A desktop computer can evaluate millions of coil configurations within seconds, enabling genuine design-space exploration that is simply infeasible with FEA.

Limitations: Analytical formulas assume idealised conditions—infinately thin conductors, absence of ferrite, no proximity effect, and perfect geometry. They cannot accurately model DD or complex hybrid coils without substantial simplification, and they neglect AC conductor losses beyond the DC resistance. AC resistance modelling for Litz wire has been addressed by Wojda and Kazimierczuk [57] and Chen et al. [13], but these models have not yet been consolidated into simple closed-form expressions applicable to arbitrary coil geometries. Lemaitre et al. [27] provided practical Litz wire design guidelines (e.g., limiting the number of bundles per stranding level to no more than five), which are incorporated into the present work.

Despite their limitations, analytical methods are frequently dismissed in the literature as insufficiently accurate, yet they retain significant value for early-stage candidate screening. The present work extends analytical methods with a per-segment proximity loss model based on Dowell’s method and discretised Neumann integration (Chapter 4), bridging the gap between simple closed-form expressions and full FEA.

3.6.3 Heuristic and Rule-of-Thumb Designs

A notable proportion of published IPT designs are based on heuristics: “use a circular coil with outer diameter 400 mm”, “place ferrite bars every 30 mm”, “use 2 mm aluminium shield”. Such rules frequently trace back to early SAE J2954 drafts or to a handful of widely cited papers [11]. Heuristics provide a rapid starting point for prototyping and require no specialised optimisation infrastructure. They offer no guarantee of optimality, however, and are inherently brittle: a rule of thumb calibrated for one power level, air gap, or misalignment envelope may perform poorly when conditions change. The prevalence of heuristic designs in the literature reflects a lack of accessible, systematic design tools. The present work aims to supplant heuristics with data-driven, automated optimisation (Chapters 4 and 5).

3.6.4 Surrogate-Assisted and Multi-Objective Optimisation

Recent work has employed surrogate models—including Kriging interpolation and neural networks—trained on FEA data to approximate the design space at a fraction of the computational cost [52, 56]. These methods can reduce the number of expensive FEA evaluations by one to two orders of magnitude. Multi-objective optimisation (Pareto front methods) is also emerging to quantify trade-offs among efficiency, mass, and cost.

Yilmaz et al. [59] applied multi-objective PSO to circular magnetic couplers, generating Pareto fronts that quantify the trade-off between coupling coefficient and leakage flux. Hassan et al. [19] employed PSO to optimise coil geometry and ferrite configuration for an LCC-compensated WPT system, achieving high DC-side efficiency (96.9% at perfect alignment and 95.5% at 150 mm misalignment), supported by sensitivity analysis. Corti et al. [15] developed MatWPT, a MATLAB toolbox combining evolutionary algorithms with Monte Carlo tolerance analysis to assess the impact of component variations on system performance. Viera et al. [50–52, 54] developed PSO-based optimisation for SS-compensated IPT chargers, including MEPT control and experimental validation at 3.7 kW.

Strengths: Surrogate and PSO methods accommodate a large number of design variables and multiple objectives simultaneously; they are systematic and reproducible.

Limitations: These methods still require a substantial upfront FEA dataset (typically tens to hundreds of simulations) or rely on analytical approximations that may not capture all relevant physical phenomena. Surrogate accuracy is highly dependent on the sampling strategy and may miss local optima in regions of sparse training data. Furthermore, the software and computational infrastructure required are not universally accessible to smaller research groups. The present work complements existing surrogate-based approaches by providing a fast analytical pre-screening stage (Chapter 4) that reduces the number of candidates requiring FEA evaluation. Unlike MatWPT, which targets system-level WPT optimisation including power electronics, this thesis focuses specifically on the magnetic coupler with a dedicated, physics-informed analytical toolchain.

3.7 Positioning of the Present Work

The preceding review reveals substantial progress in IPT magnetic coupler design alongside several persistent and consequential gaps. Most studies optimise for a single nominal alignment position, neglecting performance over the full 3-D operational envelope. Ferrite and shielding are frequently treated as independent design steps, leading to suboptimal system-level performance. Analytical modelling tools are either too simplistic to capture proximity and ferrite effects or too computationally expensive for systematic design-space exploration.

Within the same research group, the master’s theses of Pascoal [38, 39] and Viera [50, 53, 54] established a solid foundation: analytical inductance calculations were developed, coil designs were validated by FEA, and experimental prototypes were constructed. However, those analytical tools were restricted to basic circular and rectangular geometries, did not include spatial averaging

over the misalignment envelope, and treated ferrite and shielding as separate, sequential FEA steps rather than co-optimised design variables. More recently, Hassan et al. [19] applied PSO to coil and ferrite geometry for an LCC-compensated WPT system, achieving high efficiency with supporting sensitivity analysis. Although methodologically similar to the present thesis, their work relies on FEA for each PSO evaluation and does not address spatial averaging. Corti et al. [15] developed MatWPT for system-level optimisation with tolerance analysis, but the toolbox targets the complete WPT system including power electronics, rather than the magnetic coupler in isolation.

Building on these prior efforts, this thesis introduces a multi-stage methodology that fills the identified gaps through the following key contributions:

1. **Fast analytical toolchain with spatial averaging** (Chapter 4): A toolchain based on the discretised Neumann integral and a per-segment proximity loss model computes mutual inductance, self-inductance, and AC resistance for arbitrary planar coil geometries (circular, rectangular, DD, multi-layer). By evaluating hundreds of misalignment positions in seconds, it enables the computation of spatial statistics (mean, minimum, coefficient of variation) of k over the full SAE J2954 Z1 envelope. This directly addresses the single-position optimisation gap that pervades the prior literature, including earlier works. Accuracy of the analytical toolchain is established through validation against FEA simulations.
2. **Structured ferrite and shielding integration** (Chapter 5): Rather than embedding ferrite parameters into an optimisation loop (requiring FEA at every evaluation) the present work employs a staged approach in which the coil mechanical designs identified by the analytical toolchain and validated by FEA are taken as fixed inputs. A defined set of ferrite tile layouts, each combined with a fixed aluminium backplate for passive shielding, is then swept in Ansys Maxwell. Layout selection is governed by a cost-performance metric based on the maximum expected efficiency of the SS-compensated system at resonance, evaluated over the first-quadrant misalignment region. This is fundamentally different from the single-point efficiency assessment prevalent in the literature, and it ensures that the chosen ferrite configuration is matched to the spatial coupling characteristics of the coil it accompanies rather than optimised solely for aligned operation.
3. **Experimental validation of the design methodology** (Chapter 6): A prototype coupler is fabricated for a selected candidate design, encompassing both air-core and fully assembled configurations, and characterised for self-inductance, mutual inductance, and coupling coefficient across a range of separations and lateral displacements. The primary objective of this experimental campaign is to validate the physical implementation against FEA predictions from the co-design stage, thereby confirming that the complete design pipeline – from analytical pre-screening through FEA refinement to hardware realisation – yields couplers whose measured performance matches the simulated behaviour.

This thesis is concerned exclusively with the design and characterisation of the magnetic coupler; the power electronics and the broader system-level control architecture are deliberately

outside its scope. Throughout the analytical toolchain and the subsequent FEA-based pipeline, the SS compensation topology is assumed, with an ideal maximum-efficiency tracking mechanism capable of impedance matching at resonance. The classification metric used throughout this work to compare and select candidate designs, namely the maximum coil-to-coil efficiency, is itself derived under these assumptions: ideal resonance and ideal impedance matching for the SS topology (introduced in Section 2.6). Extending the methodology to alternative compensation topologies, or relaxing the ideal-tracking assumption to account for the dynamics and losses of a practical control implementation, would require corresponding modification of the efficiency expressions underpinning the toolchain, and is left as a direction for future work. Similarly, the numerical specifications used to instantiate and validate the design pipeline – coil mechanical dimensions, conductor constructions and the evaluation grid over the SAE J2954 Z1 envelope – are not imposed by external project requirements, which do not yet fix these parameters, but are adopted as a self-consistent representative framework to demonstrate the methodology. The rationale for these choices is detailed in Chapter 5. Further restrictions and assumptions specific to each stage of the pipeline are discussed, with greater rigour and in their proper context, in the corresponding chapters.

In summary, this thesis does not aim to replace FEA or experimental prototyping. Rather, it provides a systematic, efficient, and transparent methodology for navigating the large design space of IPT magnetic couplers, ensuring that computationally expensive FEA evaluations and experimental resources are applied exclusively to genuinely promising candidate designs.

Chapter 4

Analytical Toolchain and Validation

4.1 Rationale for the Toolchain

Finite element analysis provides high-fidelity solutions to Maxwell’s equations and is widely regarded as the gold standard for IPT magnetic coupler design. However, as discussed in Section 3.6, FEA is computationally expensive: a single 3D simulation can take minutes to hours, making it impractical for exhaustive parametric exploration or evolutionary optimisation. Closed-form analytical formulas (e.g., Wheeler’s formula) are fast but limited to simple geometries and cannot predict proximity losses reliably. Most practical studies also evaluate coupling only at a single nominal alignment position, which does not reflect the full SAE J2954 misalignment envelope (Section 3.4).

The analytical toolchain presented in this chapter is intended as the air-core pre-design stage of the methodology. It computes self-inductance and mutual inductance using the discretised Neumann integral (Eq. 2.12) and estimates AC resistance using a per-segment Biot–Savart proximity loss model (Section 2.3). The toolchain is deliberately restricted to air-core coils: ferrite materials and electromagnetic shielding are not modelled at this stage. This simplification is required for rapid design ranking while maintaining sufficient accuracy for pre-design screening. Ferrite and shielding are added later in Stage 3 of the five-stage pipeline (Chapter 5).

The toolchain comprises four standalone tools, each addressing a specific design task. They are implemented as MATLAB classes (`CoilClass`, `CouplingClass`) and driver scripts (`Tool_1_Litz_Wire.m` to `Tool_4_Coupling_Optimization.m`). The following sections describe each tool, its mathematical foundation, and typical use cases. Section 4.6 validates that the toolchain predicts the fundamental quantities L , M , and R (R_{dc} and R_{ac}) with sufficient accuracy for design ranking and optimisation.

4.2 Tool 1 – Litz Wire Dimensioning

Tool 1 assists the designer in specifying a Litz wire that meets the current-carrying requirements while minimising high-frequency losses. The user provides the RMS current, the operating

frequency range, the maximum allowable current density, and the operating temperature. The tool then performs the skin-depth calculation, strand diameter selection, AWG rounding, strand count determination, and bundle diameter estimation.

4.2.1 Mathematical Foundation

The skin depth at the maximum frequency is given by Eq. 2.5 from Chapter 2. To ensure negligible skin effect within each strand, the strand diameter is chosen as a small multiple of the skin depth: $d_{\text{strand}} = c_{\delta} \cdot \delta$, where the safety factor c_{δ} is user-defined (typically between 1.0 and 1.2). The nearest American Wire Gauge (AWG) number is selected by rounding up to the next available gauge, because not all AWG sizes are commercially manufactured. An additional fineness step can be specified by the user to further increase the strand count if desired.

The effective copper cross-section of a stranded wire is reduced by the enamel insulation coating on each strand. This reduction is accounted for implicitly by rounding up the required number of strands, as the tool uses the nominal copper diameter (without insulation) for the strand cross-section. The required total copper area (excluding insulation) is determined from the RMS current and the maximum allowable current density (Eq. 4.1):

$$A_{\text{copper}} = \frac{I_{\text{rms}}}{J_{\text{max}}}. \quad (4.1)$$

The number of strands is then (Eq. 4.2):

$$n_{\text{strands}} = \left\lceil \frac{A_{\text{copper}}}{\pi(d_{\text{strand}}/2)^2} \right\rceil. \quad (4.2)$$

The ceiling operation ensures that the total copper area is sufficient, and the rounding-up factor compensates for the enamel insulation and any manufacturing tolerances. This rounding behaviour is intentionally conservative: rather than attempting to predict the exact commercial Litz wire construction, Tool 1 establishes a lower bound on strand count that guarantees the required current-carrying capacity. The final selection of a specific, commercially available construction remains a practical decision for the designer, who must weigh supplier availability against the resulting current density.

The bundle outer diameter is estimated from the strand diameter, the number of strands, and an assumed packing factor p_f (Eq. 4.3):

$$d_{\text{bundle}} = d_{\text{strand}} \sqrt{\frac{n_{\text{strands}}}{p_f}}. \quad (4.3)$$

If a measured bundle diameter is known (e.g., from a datasheet), the user may supply it and the tool back-calculates the actual packing factor.

The DC resistance per metre is computed using Eq. 2.4 with the total copper area $A_{\text{total}} = n_{\text{strands}} \cdot \pi(d_{\text{strand}}/2)^2$. The tool outputs a `conductor` structure containing the bundle diameter, strand diameter, strand count, resistivity, and the computed DC resistance per metre and skin factor.

4.2.2 Typical Use and Example

Tool 1 is typically run once at the beginning of a design project, before any coil geometry is defined. The resulting `conductor` structure is saved and reused across Tools 2, 3, and 4. The designer may iterate between Tool 1 and later tools if the initial current estimate proves inaccurate.

Tables 4.1 and 4.2 show a representative set of user inputs and the corresponding outputs produced by the tool. The values are illustrative; actual outputs depend on the specific application requirements.

Table 4.1: User inputs for Tool 1.

Parameter	Value	Unit
Min frequency $f_{\min}$	80	kHz
Max frequency $f_{\max}$	90	kHz
Temperature T_{op}	85	°C
RMS current I_{rms}	30	A
Current density $J_{\max}$	3.00	A mm ⁻²
Safety factor c_{δ}	1.0	–
AWG round-up	1	–
Bundle mode	fixed	–
Bundle diameter	4.60	mm

Table 4.2: Outputs for Tool 1.

Quantity	Value	Unit
Skin depth δ	0.247	mm
Skin factor F_{skin}	0.0008	–
Strand AWG	AWG 32	–
Strand diameter	0.202	mm
Number of strands	313	–
Current per strand	0.096	A
Current density	2.99	A mm ⁻²
Bundle diameter	4.600	mm
Packing factor	0.603	–
Total copper area	10.02	mm ²
DC resistance	2.159	mΩ m ⁻¹

The tool computes the strand count as the ceiling of the theoretical value, which may result in a non-standard number (e.g., 313 strands of AWG 32). The designer must verify whether the computed construction is available from suppliers. If a standard construction is preferred, the user may simply round the strand count down to the nearest standard value (e.g., from 313 to 300 strands) and re-run the tool with an adjusted current density limit, or manually edit the output `conductor` structure. In the example above, reducing the strand count to 300 would increase the current density from 2.99 to 3.12 A/mm², which is often acceptable. The tool is intentionally conservative, allowing the designer to make the final decision based on practical availability.

4.3 Tool 2 – Single Coil Analysis

Tool 2 analyses a single coil (transmitter or receiver) in isolation. The user provides the coil geometry (outer length, width, corner radius, inter-turn spacing, number of turns, and coil topology: 'spiral' or 'DD'), the conductor specification from Tool 1, the operating frequency (or a sweep range), the operating temperature, and for multi-layer coils, the number of layers and the layer spacing.

4.3.1 Geometry Discretisation and Conductor Length

The coil path is generated turn by turn from the outer perimeter inward. For spiral coils, the path follows a rounded rectangular (or circular) contour; for DD coils, two mirrored spiral sub-coils are placed side by side. Each turn is discretised into straight segments; the number of segments per quarter-turn is user-selectable ('fast': 10, 'normal': 25, 'precise': 60). Finer discretisation improves accuracy at the cost of longer computation time. The total conductor length is the sum of the Euclidean norms of all segment vectors.

4.3.2 DC and AC Resistance

The DC resistance at the operating temperature follows Eq. 2.4 with the temperature-corrected resistivity $\rho(T)$. The AC resistance is modelled as the incremental loss term $R_{ac} = R_{dc}(F_{skin} + F_{prox})$, such that the total winding resistance is $R = R_{dc} + R_{ac}$.

4.3.2.1 Skin Factor F_{skin}

The skin effect within each strand is modelled using Ferreira's approximation for round conductors. Let r_s be the strand radius and δ the skin depth (Eq. 2.5). The ratio $\xi = r_s/\delta$ determines the severity of skin effect. Ferreira's formula provides a smooth transition between low and high frequencies (Eq. 4.4):

$$F_{skin} = (1-t) \cdot \frac{\xi^4}{48} + t \cdot \left(\frac{\xi}{\tanh \xi} - 1 \right), \quad (4.4)$$

where the weighting function $t = 1/(1 + e^{-5(\xi-1.5)})$ ensures a gradual transition from the low-frequency approximation ($\xi^4/48$) to the high-frequency approximation ($\xi/\tanh \xi - 1$). This expression is accurate for solid round wires and, when applied to each strand of a Litz wire, accurately captures the skin effect within the individual strands.

4.3.2.2 Proximity Factor F_{prox}

Proximity losses arise from the external magnetic fields produced by neighbouring turns and other layers. The tool computes the tangential RMS magnetic field H_{rms} at the centre of each segment using the Biot–Savart law, summing contributions from all other segments in the same coil (including all layers). For a Litz wire, the proximity loss density per unit length is given by (Eq. 4.5):

$$p_{prox} = \frac{\pi^2 n_s \sigma \mu_0^2 \omega^2 d_s^4}{128} H_{rms}^2, \quad (4.5)$$

where n_s is the number of strands, σ the conductivity, $\omega = 2\pi f$, and d_s the strand diameter. This expression is derived from the exact solution for a round conducting wire in a uniform transverse alternating magnetic field, as adapted by Ferreira for Litz wire bundles. The total proximity loss is

obtained by summing $p_{\text{prox}} \cdot l_i$ over all segments of length l_i . The proximity factor is then:

$$F_{\text{prox}} = \frac{\sum_i p_{\text{prox},i} l_i}{R_{\text{dc}} I^2}. \quad (4.6)$$

Because the field is evaluated at the centre of each segment, the method accounts for variations in field magnitude along the coil (e.g., stronger fields near the centre of the winding and weaker fields near the periphery). For multi-layer coils, contributions from all layers are included, capturing inter-layer proximity effects. This per-segment Biot–Savart evaluation is a key methodological advantage of the toolchain: unlike standard one-dimensional models (such as Dowell’s method), which assume a uniform field across the entire winding cross-section, this approach captures spatial field gradients along the coil’s own geometry. Since the magnetic field experienced by a turn near the centre of a spiral winding differs substantially from that experienced by a turn at the periphery, this segment-resolved treatment yields a markedly more accurate proximity loss estimate for the non-uniform field distributions characteristic of planar IPT coils.

4.3.3 Self-Inductance and Quality Factor

For a single-layer coil, self-inductance is computed via the discretised Neumann integral (Eq. 2.12), with the self-term replaced by the logarithmic approximation for a round wire (Section 2.4). For multi-layer coils, the total self-inductance is the sum of the self-inductances of each layer plus twice the sum of the mutual inductances between all layer pairs. The mutual inductance between two layers is computed using the same Neumann integral, accounting for the vertical offset between their midpoints. The inductance is assumed frequency-independent, which is valid for air-core coils up to several hundred kilohertz.

The quality factor is then $Q = \omega L/R$. Tool 2 can perform a frequency sweep, storing R_{dc} (constant), $R_{\text{ac}}(f)$, $R(f)$, L (constant), and $Q(f)$. Results can be exported to CSV for further analysis.

4.3.4 Typical Use and Example

Tool 2 is used to characterise a single coil across a set of frequencies before it is placed in a transmitter–receiver pair. Typical use cases include validating a candidate design before coupling analysis, comparing different Litz wires or turn counts, and performing sensitivity studies on spacing or corner radius.

Figure 4.1 shows a 3D visualisation of a double-layer rectangular spiral coil with the parameters listed in Table 4.3. The conductor used for this example is the same as defined in Table 4.2. The computed outputs are given in Table 4.4.




Figure 4.1: 3D views of the example spiral coil: (a) angled view, (b) top view.

4.4 Tool 3 – Spatial Coupling Analysis

Tool 3 evaluates the magnetic coupling between a transmitter coil (Tx) and a receiver coil (Rx) over a three-dimensional spatial region. The user supplies the Tx and Rx coil definitions (each can be any geometry supported by Tool 2), the evaluation region (ranges and step sizes for x , y , and the air gap z), and optionally the operating frequency and temperature.

4.4.1 Spatial Sweep and Coupling Metrics

For each grid point (x, y, z_{gap}) , the mutual inductance M is computed using the Neumann integral between Tx and Rx (Eq. 2.12). The centre-to-centre distance accounts for the physical stack heights of multi-layer coils: if a coil has multiple layers, its face height (distance from the stack centre to the outermost surface) is added to the air gap to obtain the centre-to-centre separation used in the Neumann integral. The coupling coefficient then follows from Eq. 2.2.

The maximum achievable coil-to-coil efficiency, assuming ideal resonant compensation and optimal load, is given by Eq. 2.24, which, for a given fixed topology, justifies decoupling the design of the magnetic coupler from the design of the power electronics.

4.4.2 Exploitation of Symmetry and Grid Discretisation

The user defines the evaluation grid by specifying ranges and step sizes for x , y , and z . For coil geometries that are symmetric with respect to the coordinate axes, the mutual inductance $M(x, y)$ and thus all derived quantities inherit that symmetry. Consequently, the user may restrict the evaluation region to the first quadrant ($x \geq 0$, $y \geq 0$) without loss of information; the toolchain automatically mirrors the results when calculating statistics over the full envelope.

The symmetry argument is geometric in nature. Circular coils are rotationally symmetric about their centre, and square coils centred at the origin are symmetric about both the x and y axes. Rectangular coils centred at the origin are also symmetric about both axes, even when side lengths differ (see Section 3.2). In all these cases, the coil footprint is identical when reflected across either

Table 4.3: Example inputs for Tool 2.

Parameter	Value	Unit
Topology	spiral	–
Length	500	mm
Width	400	mm
Corner radius	100	mm
Turn count	10	–
Turn spacing	2	mm
Layer count	2	–
Layer spacing	2	mm
Precision level	normal	–
Frequency F_{op}	85	kHz
Temperature T_{op}	25	°C

Table 4.4: Example outputs for Tool 2.

Quantity	Value	Unit
R_{dc}	52.67	mΩ
R_{ac}	171.44	mΩ
R	224.11	mΩ
L	285.35	μH
Q	680.00	–
Length	28.80	m

axis, and since the transmitter and receiver are co-centred in the nominal alignment, the coupling coefficient satisfies $k(x, y) = k(-x, y) = k(x, -y) = k(-x, -y)$, so full quadrant symmetry applies.

Double-D coils, as noted in Section 3.2, possess only half-plane symmetry: they are symmetric about the axis perpendicular to the winding split (typically the y -axis) but not about the split axis itself. For a DD coil split along the x -direction, $k(x, y) = k(-x, y)$; symmetry about the x -axis is not guaranteed. In that case, the user may restrict the region to $x \geq 0$ over the full y range (half-plane). This distinction is particularly important for DD–DD coupling pairs, where the symmetry-reduction logic of Table 4.5 cannot be applied directly in its quadrant form and must instead use the half-plane reduction to avoid discarding genuinely asymmetric information.

For coils with full quadrant symmetry, the user can therefore define the evaluation region using only positive ranges (e.g., $x \in [0, 75]$ mm, $y \in [0, 100]$ mm) instead of symmetric negative ranges. Table 4.5 shows the resulting reduction in the number of evaluation points for the SAE J2954 Class ZI envelope (with a step size of 25 mm). The full symmetric region would require $7 \times 9 \times 3 = 189$ points; restricting to the first quadrant reduces this to $4 \times 5 \times 3 = 60$ points – a saving of more than a factor of three.

Table 4.5: Number of evaluation points for the SAE J2954 Class ZI envelope (step size 25 mm).

Region defined by user	Total points
Full region (x and y from negative to positive)	$7 \times 9 \times 3 = 189$
First quadrant only ($x \geq 0, y \geq 0$)	$4 \times 5 \times 3 = 60$

The choice of step size involves a trade-off between computational effort and spatial resolution. A finer step (e.g., 10 mm) increases the number of points, improving the accuracy of spatial statistics but prolonging computation. A coarser step (e.g., 50 mm) reduces the number of points and speeds up the analysis, but may smooth over local variations in coupling. Two different coil designs that have comparable average efficiency on a coarse grid may diverge significantly when evaluated on a finer grid, because the coarse grid could miss narrow regions of poor coupling. Crucially, the step

size also determines the cost of Stage 2 FEA validation, since every analytically evaluated grid point that advances to validation requires a separate Ansys Maxwell simulation; halving the step size would roughly quadruple the number of required FEA runs for the same spatial extent. In this work, a step size of 25 mm was adopted as the optimal compromise between the statistical fidelity of the spatial sweep and the computational cost it imposes downstream on FEA validation, and is consistent with typical grid resolutions used in the literature for SAE J2954 compliance tests.

The symmetry reduction is particularly valuable because the same spatial grid is later simulated in Ansys Maxwell, where each position requires a separate 3D FEA setup. Thus, the symmetry-aware definition of the evaluation region not only accelerates the analytical pre-design but also significantly reduces the number of costly FEA simulations needed for validation.

Figure 4.2 illustrates the reduction: the full region (left) with extents ± 75 mm in x and ± 100 mm in y contains four symmetric quadrants; the right figure shows the single quadrant ($0 \leq x \leq 75$ mm, $0 \leq y \leq 100$ mm) that is actually evaluated when the user defines the region accordingly. The red points mark the evaluated positions of the receiver coil centre; the blue point is the fixed transmitter centre.




Figure 4.2: Point cloud of receiver coil centres (red) for a spatial analysis: (a) full region (189 points); (b) first quadrant (60 points). The transmitter coil centre is fixed at the origin (blue).

4.4.3 Statistics and Performance Optimisation

After the sweep, Tool 3 computes the minimum, maximum, and average values of M , k , and $\eta_{\max}$ over the entire defined region. These statistics directly address the literature gap identified in Section 3.4, where most studies evaluate coupling only at a single nominal alignment position. The full three-dimensional arrays are stored for optional plotting, and the data can be exported to CSV.

To achieve acceptable runtime, the tool employs several optimisations: the Rx geometry is extracted once and only its midpoint coordinates are shifted per grid point; the Tx sub-coils are decomposed into plain structs that can be efficiently broadcast in a `parfor` parallel loop; and the

Neumann integral uses a blocking strategy to bound memory usage. Tool 3 is the workhorse for comparing competing coil designs under realistic misalignment conditions.

4.4.4 Typical Use and Example

The following example illustrates a typical use of Tool 3. A two-layer spiral transmitter (500 mm $\times$ 400 mm) and a single-layer spiral receiver (300 mm $\times$ 300 mm) are analysed. The evaluation region is defined in Table 4.6, and the complete input parameters for both coils are listed in Table 4.7. The conductor used is the same as defined in Table 4.2.

Table 4.6: Evaluation region for Tool 3 example.

Axis	Range (min : step : max)	Unit
x (lateral)	0 : 25 : 75	mm
y (lateral)	0 : 25 : 100	mm
z (air gap)	100 : 25 : 150	mm

Table 4.7: Input parameters for Tx and Rx coils.

Parameter	Tx	Rx	Unit
Topology	spiral	spiral	–
Length	500	300	mm
Width	400	300	mm
Corner radius	100	50	mm
Turn count	10	8	–
Turn spacing	2.0	2.0	mm
Layer count	2	1	–
Layer spacing	2.0	–	mm
Precision level	normal		–
Frequency f_{op}	85		kHz
Temperature T_{op}	25		°C

Figure 4.3 shows a 3D visualisation of the Tx and Rx coils at a specific offset ($x = 50$ mm, $y = 50$ mm, $z_{\text{gap}} = 150$ mm).

The coupling statistics computed over the entire evaluation region are summarised in Table 4.8, which reports the minimum, maximum, and average values of mutual inductance M , coupling coefficient k , and maximum efficiency η_{max} .

In addition, Tool 3 outputs a summary of the electrical parameters for each coil in the same format as Tool 2.




Figure 4.3: 3D visualisation of Tool 3's Tx (blue) and Rx (red) coils at a given offset (50 mm lateral in x and y, 150 mm air gap).

4.5 Tool 4 – Coupling Optimisation

Tool 4 performs global optimisation of the Tx and Rx geometries to maximise a user-defined coupling metric over the spatial region. The user provides the search space bounds for each geometric parameter, discrete step sizes (to enforce manufacturable dimensions), the objective function (choice of metric and statistic), optional electrical constraints, a symmetry flag (whether Tx and Rx are forced identical), and PSO hyperparameters.

4.5.1 Why Optimisation Is Needed

As illustrated in Tool 3, the magnetic coupling between Tx and Rx depends on a large number of geometric parameters: outer length, outer width, corner radius, turn spacing, number of turns, layer count, layer spacing, and topology (spiral or DD orientation). For each coil, up to eight independent parameters can be varied. If both Tx and Rx are optimised independently, the total number of free parameters can reach sixteen. A brute-force parametric sweep over such a high-dimensional

Table 4.8: Coupling statistics for Tool 3 example.

Quantity	Min	Max	Avg
M (μH)	10.09	19.77	14.78
k	0.1116	0.2186	0.1634
$\eta_{\max}$ (%)	97.13	98.52	97.97

space is computationally prohibitive, even with the fast analytical toolchain. Moreover, the spatial coupling metrics (minimum, average, or maximum efficiency over the SAE J2954 region) are non-linear and may exhibit multiple local optima. Therefore, a global optimisation algorithm is required.

4.5.2 Particle Swarm Optimisation (PSO) – Overview

Particle Swarm Optimisation is a population-based stochastic search method inspired by the collective behaviour of bird flocks or fish schools. A set of candidate solutions (particles) moves through the search space, each particle updating its position based on its own best known position (personal best) and the best position found by the entire swarm (global best). The velocity update for each particle is given by (Eq. 4.7):

$$\mathbf{v}_i(t+1) = w\mathbf{v}_i(t) + c_1r_1(\mathbf{p}_i - \mathbf{x}_i(t)) + c_2r_2(\mathbf{g} - \mathbf{x}_i(t)), \quad (4.7)$$

where w is the inertia weight, c_1 and c_2 are acceleration coefficients, and r_1, r_2 are random numbers uniformly distributed in $[0, 1]$. The position is then updated as $\mathbf{x}_i(t+1) = \mathbf{x}_i(t) + \mathbf{v}_i(t+1)$.

PSO requires neither differentiability of the objective function nor convexity. It is well-suited for the IPT coupler optimisation problem because the design space is moderately sized (8–16 variables), the objective function (efficiency over a region) is expensive to evaluate (each evaluation requires a full 3D coupling sweep), and the presence of constraints (geometric and electrical) can be handled naturally through infeasibility penalties.

Figure 4.4 illustrates the concept: initially, particles are randomly distributed across the search space; after several iterations, they converge toward the region of the global optimum.




Figure 4.4: Illustration of particle swarm evolution: (a) initial random positions, (b) convergence starting after several iterations.

4.5.3 Parameter Space and Discretisation

The user defines lower and upper bounds for each optimisable parameter. To fix a parameter the lower and upper bounds are set to the same value. Before evaluating the objective, each candidate is snapped to the nearest discrete grid point defined by user-supplied step sizes. This ensures that the final design is manufacturable. Integer parameters (turns, layer count) are rounded to the nearest integer after snapping. This snap-to-grid discretisation is what distinguishes the present formulation from a purely continuous optimisation: without it, PSO would readily converge on geometrically optimal but practically unrealisable dimensions (e.g., a turn spacing of 1.73 mm or a fractional turn count). By constraining every candidate evaluated by the swarm to lie on a grid of physically realisable values, the algorithm is guaranteed to return a design that can be directly manufactured, without requiring a separate post-hoc rounding step that could invalidate the optimality of the result.

The topology parameter is treated as a discrete variable with three possible values: 0 (spiral), 1 (DD-x, split along X), 2 (DD-y, split along Y). The PSO algorithm handles this by treating it as a continuous variable in $[0, 2]$ and rounding to the nearest integer before constructing the coil.

4.5.4 Objective Function and Constraints

The user selects a coupling metric (mutual inductance M ('M'), coupling coefficient k ('K'), or maximum efficiency $\eta_{\max}$ ('N')) and a statistic (average ('avg'), minimum ('min'), maximum ('max'), or standard deviation ('std')). The objective function is defined as the negative of the chosen statistic because PSO minimises. For instance, to maximise the average efficiency over the defined region, the objective is (Eq. 4.8):

$$\Phi = -\frac{1}{N_{\text{points}}} \sum_{i=1}^{N_{\text{points}}} \eta_{\max}(x_i, y_i, z_i). \quad (4.8)$$

If the user selects 'min' instead, the objective becomes the negative of the worst-case efficiency, promoting designs that are robust to misalignment.

Constraints are implemented as penalty-free infeasibility: if any geometric constraint (e.g., inner diameter becomes negative, corner radius exceeds half the minimum side) or electrical constraint (e.g., inductance outside a user-specified range) is violated, the objective function returns a large positive value (10^9). PSO naturally avoids such infeasible regions because they have poor objective values. Electrical constraints are useful when the power electronics impose limits on coil inductance or resistance; for example, the compensation network may require L to stay below a threshold to avoid excessively high voltages. The user supplies such limits via the `elec_limits` structure (`R`, `L`, `Q`). These limits apply to the air-core values; the addition of ferrite and shielding will modify them, but the user can re-run the optimisation with adjusted constraints if needed.

4.5.5 Typical Use and Example

The following example demonstrates a typical optimisation scenario: the transmitter coil is fixed (with the same parameters as in Tools 2–3), while the receiver coil is optimised. The receiver search space is presented in Table 4.9. The conductor is the same as used in Tools 1–3 (Table 4.2). The evaluation region is identical to that in Tool 3 (Table 4.6). The objective is to maximise the average efficiency $\eta_{\max}$ over the region.

The optimisation was run with a swarm size of 20 particles for 50 iterations – each iteration assesses 20 candidate designs, where each candidate requires a complete 3D spatial sweep over the 60 grid points in the SAE J2954 region (20 particles $\times$ 60 spatial points = 1,200 coil evaluations per iteration).

Table 4.9: Input coil search space for Tool4 example.

Parameter	Tx		Rx		Unit
	Min	Max	Min	Max	
Topology	spiral		any		–
Length	500	500	100	300	mm
Width	400	400	100	300	mm
Corner radius	100	100	0	150	mm
Turn count	10	10	2	20	–
Turn spacing	2.0	2.0	0	10	mm
Layer count	2	2	1	1	–
Layer spacing	2.0	2.0	–	–	mm
Free parameters count	0		6		–
Precision level	normal				–
Frequency f_{op}	85				kHz
Temperature T_{op}	25				°C

After optimisation, the optimal Rx parameters are obtained. Table 4.10 lists the resulting geometry, and Table 4.11 gives the coupling statistics over the evaluation region. Figure 4.5 shows a 3D visualisation of the fixed Tx and the optimised Rx at the nominal position.

For the presented example we can compare the arbitrated Rx design from Tool 3 with the optimised design from Tool 4. The average efficiency over the region improved from 97.97% to 98.05%, which may seem small, but it shows how modest parameter changes can improve the maximum achievable efficiency. The average mutual inductance M decreased from 14.78 μH to 13.95 μH , implying weaker coupling. However, the optimised Rx coil also reduced its total resistance R from 27.55 m Ω to 22.59 m Ω , which lowers losses and compensates for the loss in coupling. This example illustrates the complex interplay between geometry, coupling, and losses, and how the optimisation tool can navigate this trade-off to find improved designs.

The optimised receiver geometry can then be exported for further analysis (FEA with ferrite and shielding, thermal simulation, etc.). The toolchain allows the user to iterate: if the electrical constraints are not satisfied after adding ferrite, the optimisation can be re-run with tightened limits.

Table 4.10: Optimised Rx parameters for Tool 4 example.

Parameter	Value	Unit
Topology	spiral	–
Length	300	mm
Width	300	mm
Corner radius	85	mm
Turn count	9	–
Turn spacing	4.0	mm
Layer count	1	–
Layer spacing	–	–

Table 4.11: Coupling statistics for Tool 4 example.

Quantity	Min	Max	Avg
M (μH)	9.57	18.47	13.95
k	0.1101	0.2126	0.1606
$\eta_{\max}$ (%)	97.25	98.57	98.05

4.6 Validation of Air-Core Predictions

Before the toolchain can be deployed for design iteration, its predictions must be systematically validated against high-fidelity finite-element software. This section describes the validation methodology, the comprehensive test matrix, and the observed accuracy across a wide range of design scenarios. All comparisons employ air-core configurations (no ferrite, no shielding) since this represents the scope of the analytical toolchain presented above.

4.6.1 Test Matrix Overview

A dedicated MATLAB script, `Validation_Tool.m`, generates and evaluates diverse transmitter–receiver (Tx–Rx) coupling pairs. Ten families of test cases systematically explore different design aspects. Table 4.12 summarises these families; detailed geometric and conductor parameters for each case are presented in Appendix A.

Table 4.12: Summary of test coupling configurations for validation.

Group	Transmitter (Tx)	Receiver (Rx)	Design Aspect Tested
T1	Single-layer Spiral	Single-layer Spiral	Geometric variation
T2	Single-layer Spiral	Single-layer Spiral	Size variation
T3	Single-layer Spiral	Single-layer Spiral	Conductor variation
T4	Single-layer Double-D	Single-layer Spiral	Mixed topology
T5	Single-layer Double-D	Single-layer Double-D	Both coils Double-D
T6	Multi-layer Spiral	Single-layer Spiral	Multi-layer effects
T7	Multi-layer Spiral	Multi-layer Spiral	Both coils multi-layer Spiral
T8	Multi-layer Double-D	Single-layer Spiral	Mixed topology
T9	Multi-layer Double-D	Multi-layer Double-D	Both coils multi-layer Double-D
T10	Mixed design	Mixed design	Mixed geometries and conductors

For each coupling pair, the toolchain computes fundamental electromagnetic quantities: total resistance $R_{\text{Tx}}, R_{\text{Rx}}$ (including both DC and AC components), self-inductance $L_{\text{Tx}}, L_{\text{Rx}}$, and mutual




Figure 4.5: 3D visualisation of Tool 4's Tx (blue) and Rx (red) coils at a given offset (50 mm lateral in x and y, 150 mm air gap).

inductance M . All derived metrics – quality factor Q , coupling coefficient k , and maximum efficiency $\eta_{\max}$ – are obtained from these fundamental quantities. Consequently, even small errors in the base quantities propagate to derived metrics, sometimes amplifying in magnitude.

4.6.2 Conductor Constructions

Table 4.13 lists the four Litz wire types used throughout the validation. These constructions span a practical range: from coarse wire (LITZ_A, 300 strands) suitable for low-frequency, high-current applications, to fine wire (LITZ_B and LITZ_C, 900 strands each) used at higher frequencies, to an ultra-coarse wire (LITZ_D, 100 strands) for extreme cases. This variety exercises the frequency-dependent resistance model (skin effect and proximity effect) across realistic design parameter space.

Table 4.13: Summary of Litz wire constructions used in validation test matrix.

ID	Strands	AWG	Bundle diam (mm)	Packing factor
LITZ_A	300	32	4.00	0.770
LITZ_B	900	38	6.00	0.254
LITZ_C	900	38	5.00	0.365
LITZ_D	100	28	5.00	0.425

4.6.3 Evaluation Methodology

All test cases are evaluated under two distinct parameter sweeps, each targeting a complementary aspect of the toolchain's functionality:

1. **Frequency sweep:** DC (0 Hz), 70 kHz, 85 kHz, and 100 kHz are evaluated at fixed ambient temperature of 25 °C. The multiple frequency points exercise the frequency-dependent resistance model (including skin effect and proximity effect computations) across the full range of conductor types (`LITZ_A` through `LITZ_D`). This validates the toolchain's ability to capture resistance variation with frequency.
2. **Spatial sweep:** A representative set of eight positions samples the extrema and corners of the SAE J2954 *Z1* misalignment region. Specifically:
 - Lateral misalignment: $(x, y) \in \{0, 75\} \times \{0, 100\}$ mm,
 - Vertical separation: $z \in \{100, 150\}$ mm.

This yields positions: (0, 0, 100), (75, 0, 100), (0, 100, 100), (75, 100, 100), and the corresponding four at $z = 150$ mm.

Each coupling pair thus experiences $4 \text{ frequencies} \times 8 \text{ positions} = 32$ evaluations. However, at DC ($f = 0$ Hz), mutual inductance is not physically meaningful, so only DC resistance validation applies; this leaves $3 \times 8 = 24$ AC frequency evaluations per pair. With a test matrix of T_1 through T_{10} families spanning approximately 36 distinct coupling pairs (see Appendix A), the complete validation dataset comprises:

$$36 \text{ pairs} \times 24 \text{ AC evaluations} = 864 \text{ total evaluation points}$$

This scale ensures that error statistics are robust and not dominated by individual outlier cases.

4.6.4 Comparison with Ansys Maxwell

Each test case is simulated identically in Ansys Maxwell 2025 R2, using the same coil geometry, conductor specifications, frequency points, and spatial positions. The electromagnetic quantities computed by the MATLAB toolchain ($R_{dc}, R_{ac}, L, M, k, \eta_{max}$) are extracted from Maxwell's finite-element solution. The percent error for any quantity X is defined symmetrically as:

$$\delta_X = \frac{X_{\text{MATLAB}} - X_{\text{Maxwell}}}{X_{\text{Maxwell}}} \times 100\%. \quad (4.9)$$

This choice ensures that overestimation and underestimation are equally weighted; the choice of denominator (Maxwell result) reflects the convention that FEA provides the reference high-fidelity value. For very small values of X_{Maxwell} , the percent error can become artificially large even when the absolute error is small. This occurs when the transmitter is a double-D coil and the receiver is positioned along the coil's symmetry axis, where mutual inductance approaches zero. In such

edge cases, the percent error metric is excluded from the analysis, and only the absolute error is considered.

4.6.4.1 DC and AC Resistance

Table 4.14 summarises the percent errors for resistance validation.

DC Resistance: The DC resistance error is consistently within 0.5–3.1% for typical cases, with a maximum observed error of 5.3%. This confirms that conductor length calculations and DC resistivity values are accurate, validating the geometric discretisation and material property specifications in Tool 1.

AC Resistance: For AC resistance, typical errors lie in the range 0.7–5.8%, with a maximum of 12.8%. The largest errors occur for the coarse Litz wire (`LITZ_A`) in configurations with very tight turn spacing or multi-layer arrangements with small layer spacing. In these scenarios, strand-level proximity effects become severe and localised, violating the uniform-bundle-current assumption embedded in the proximity factor model. Such coarse wire at extremely tight spacing is uncommon in practical WPT designs; tighter packing typically uses finer wire (e.g., `LITZ_B` or `LITZ_C`) to maintain accuracy. For fine wires and moderate spacing, errors fall to the lower end of the typical range (0.7–3%).

Table 4.14: Percent errors for resistance predictions, compared to Ansys Maxwell.

Quantity	Typical error	Maximum error
R_{dc}	0.5–3.1 %	5.3 %
R_{ac}	0.7–5.8 %	12.8 %

4.6.4.2 Self-Inductance and Mutual Inductance

Table 4.15 summarises the inductance validation results.

Self-Inductance L : Typical errors for self-inductance are 0.1–1.3%, with a maximum of 3.1%. The largest errors occur for very tightly wound coils (inner diameter < 30 mm), where the filamentary (thin-wire) approximation becomes less accurate because the conductor cross-section is no longer negligible relative to the turn curvature. For most practical designs, the error remains well below 2%.

Mutual Inductance M : Mutual inductance errors range typically from 0.5% to 5.5%, with a maximum of 9.8%. The error varies with misalignment: at the aligned position, errors are near the lower end of the typical range; at extreme misalignments (75 mm in x or 100 mm in y), the absolute value of M becomes small and more sensitive to modelling inaccuracies, causing the relative error to increase. Moreover, in scenarios with DD coils, when the receiver coil is placed alongside the winding seam, the mutual inductance approaches zero due to the magnetic field distribution of DD coils (Section 3.2). In these situations, and whenever the Ansys Maxwell mutual inductance result

was practically zero, the evaluation point was eliminated from the analysis to avoid artificially inflated percent errors.

Table 4.15: Percent errors for inductance predictions, compared to Ansys Maxwell.

Quantity	Typical error	Maximum error
L	0.1–1.3 %	3.1 %
M	0.5–5.5 %	9.8 %

4.6.4.3 Coupling Coefficient and Maximum Efficiency

Table 4.16 presents the validation results for derived quantities.

Coupling coefficient k : Typical errors are 4–8%, with a maximum of 11.3%. The error is constrained by the errors in L and M ; it grows with misalignment as M error dominates, but the impact on design ranking remains small.

Maximum efficiency $\eta_{\max}$: The absolute error in efficiency is remarkably small: typically below 0.5 percentage points, with a maximum of 1.5 percentage points. This is because efficiency is a weak function of resistance when the figure of merit $k^2 Q_1 Q_2$ is large (as is typical in WPT systems). Even when the resistance prediction carries 10% error, the efficiency prediction remains useful for discriminating between good and poor designs.

Table 4.16: Errors for coupling and efficiency predictions, compared to Ansys Maxwell.

Quantity	Typical error	Maximum error
k	4–8 %	11.3 %
$\eta_{\max}$	< 0.5 % absolute	1.5 % absolute

4.6.5 Validation Closure

The comprehensive validation confirms that the toolchain’s errors are consistent across geometries, conductor types, frequencies, and misalignment positions. The toolchain correctly identifies the relative ordering of designs, which is the critical capability for a pre-design filter. Figure 4.6 summarises the distribution of percent errors for the direct parameters (R , L , M) across the entire validation dataset.

Detailed discussion of the error sources and their implications for practical use is provided in the following section.

4.7 Critical Assessment and Future Directions

The analytical toolchain presented in this chapter is deliberately scoped to the air-core pre-design stage, and its predictions should be interpreted within this scope. Stage 1 of the design




Figure 4.6: Boxplots of percent errors for all direct parameters (R , L , M) across the entire validation dataset.

pipeline (Chapter 5) operates exclusively on air-core coils: ferrite flux guides and electromagnetic shielding are not modelled at this stage, and the conductor bundle is treated under a perfect-filament assumption, with self-inductance assumed independent of frequency. These simplifications are what enable the toolchain to evaluate the full SAE J2954 spatial envelope and conduct global PSO optimisation within practical computation times. The validation results of Section 4.6 confirm that these assumptions introduce no significant penalty to the toolchain’s primary task: ranking candidate designs prior to FEA refinement.

However, the toolchain’s deliberate scope introduces several limitations that affect its generalisability, the depth of the optimisation, and the range of applications it can address. The following subsections synthesise these limitations and identify concrete directions for future development.

4.7.1 Ferrite and Shielding Are Excluded

The most fundamental limitation of the toolchain is that it models coils in free space only. Ferrite flux guides and aluminium shielding are not modelled at this stage, despite their significant influence on self- and mutual inductance, as demonstrated in Chapter 5. The addition of ferrite increases both self- and mutual inductance, and the aluminium shield introduces eddy-current losses that affect the total resistance. The air-core optimum identified by the toolchain is therefore not necessarily the optimum once these materials are introduced.

Extending the toolchain to incorporate these materials is non-trivial. A full 3D interface akin to Ansys Maxwell would sacrifice the speed that makes the toolchain valuable, while a simpler parametric representation would constrain the designer to a limited set of pre-defined arrangements. Even if an analytical ferrite model were feasible, the increased complexity would risk obscuring the toolchain’s primary value: rapid, transparent geometric screening. The current staged approach –

fast analytical screening followed by targeted FEA – represents a deliberate and pragmatic trade-off between speed, accuracy, and usability.

Future direction: Integrating lumped reluctance models for ferrite tiles would allow the toolchain to estimate self- and mutual inductance enhancements without FEA. A pragmatic approach might restrict ferrite layouts to tile-based arrangements (as used in Chapter 5) rather than attempting a full 3D model, moving ferrite layout optimisation into the high-speed Stage 1 loop for simultaneous optimisation of coil geometry, Litz wire construction, and ferrite layout. This would significantly reduce the number of FEA evaluations required in Stage 3 of the pipeline.

4.7.2 No Formal Optimality Guarantee

Tool 4 employs particle swarm optimisation to identify high-performing geometries, but there is no mathematical guarantee of global optimality – even for the air-core case. The optimisation is heuristic, and solution quality depends on PSO hyperparameters (inertia weight, acceleration coefficients), the discretisation grid, and the number of iterations. Furthermore, the air-core optimum is not necessarily the optimum once ferrite and shielding are introduced. The toolchain is explicitly positioned as a pre-design filter, not as a substitute for final FEA-based refinement.

This limitation has practical implications; a designer cannot be certain that the geometry returned by Tool 4 is the best possible within the search space, nor that it will remain optimal after ferrite integration. The toolchain provides a good starting point, but it does not replace the need for FEA validation and, if necessary, re-optimisation with ferrite included.

Future work could extend the current PSO formulation to generate Pareto fronts mapping trade-offs between ferrite material cost, pad weight, spatial efficiency, and thermal margin. Additionally, the optimisation could be run with different PSO hyperparameters or multiple independent runs to assess the sensitivity of the solution to the algorithm's stochastic nature.

4.7.3 Limited Multi-Receiver Optimisation

The current PSO formulation optimises for a specific transmitter–receiver pair. It cannot optimise a single transmitter that must perform well across diverse receiver geometries – a scenario arising in public charging infrastructure where one ground pad serves multiple vehicle models. This limitation is particularly significant for applications where interoperability is a requirement, such as SAE J2954 Class I systems.

The toolchain can evaluate a given Tx–Rx pair, but it cannot currently aggregate performance across multiple receivers within a single optimisation run. This means that a transmitter optimised for a specific receiver may perform poorly when paired with a different receiver, and the designer has no systematic way to trade off performance across multiple receiver families.

Future direction: The optimisation framework could be extended to handle multi-receiver scenarios through two complementary approaches: (i) a user-defined set of receiver geometries, where the transmitter is optimised for aggregate performance (e.g., maximising the minimum or average efficiency across all receivers); and (ii) an automated approach where the toolchain

discretises the receiver search space – selecting a representative set of receiver geometries from the user-specified parameter ranges – and optimises the transmitter for aggregate performance across this discretised set. This would produce a transmitter design that is closer to being truly universal with respect to the intended receiver families.

4.7.4 No Angular Misalignment

The toolchain assumes perfect planar alignment and ignores pitch, roll, and yaw. These degrees of freedom affect coupling and derived metrics, and their exclusion limits the toolchain's applicability to scenarios where angular misalignment is significant. In automotive applications, angular misalignment can arise from vehicle suspension variations, uneven ground surfaces, or imperfect parking. The coupling degradation due to angular misalignment can be comparable to that due to lateral displacement, and ignoring it may lead to optimistic efficiency predictions.

The Neumann integral formulation used in the toolchain could, in principle, be extended to include rotational degrees of freedom. However, this would significantly increase the computational cost of each evaluation, as the relative orientation of the coils would need to be recomputed for each grid point. The current implementation prioritises speed over completeness in this respect.

Future direction: Modifying the Neumann integral geometry to include pitch and roll degrees of freedom would enable the toolchain to capture the full six-degree-of-freedom misalignment problem. This would require a more complex geometric transformation of the coil coordinates, but it would significantly broaden the toolchain's applicability to real-world scenarios where angular misalignment is unavoidable.

4.7.5 Finite Coil Geometry Support

The toolchain supports spiral coils (circular, square, rectangular, rounded-corner) and DD coils with multi-layer windings, but not polynomial coils, hybrid shapes, or stacked multi-coil arrangements. This restricts the design space accessible to the analytical pre-filter. While the supported geometries cover the vast majority of practical IPT coupler designs, there are emerging topologies – such as polynomial coils with optimised field distributions, hybrid shapes that combine circular and rectangular features, and stacked multi-coil arrangements – that cannot be evaluated with the current toolchain.

The limitation is primarily a matter of implementation complexity. Extending the geometry discretisation to support arbitrary N-sided polygons or hybrid shapes would require significant modifications to the coil generation code and to the Neumann integral evaluation, but there is no fundamental theoretical barrier.

Future direction: Generalising the geometry discretisation and Neumann integral evaluation to support arbitrary N-sided polygons, hybrid shapes (e.g., rectangular with a circular inset), and stacked multi-coil configurations (e.g., double-layer DD with rotated layers) would significantly broaden the design space accessible to the analytical pre-filter. This would enable the toolchain

to be applied to emerging coil topologies and to compare them with conventional designs on a consistent basis.

4.7.6 Lack of Parameter Sensitivity Analysis

PSO identifies high-performing geometries but does not reveal which geometrical parameters most strongly influence L , M , R , or $\eta_{\max}$. Such analysis would provide valuable insight into the relationship between design variables and electrical performance, helping to identify which parameters are most critical for achieving high efficiency and robust coupling.

The absence of sensitivity analysis also means that the robustness of the optimised designs to variations is unknown. A geometry that sits on a narrow ridge of the design space may be optimal in theory but impractical in production or sensitive to small changes in dimensions, while a slightly less optimal geometry that is more tolerant to variations may be preferable in practice. Understanding these trade-offs would enable more informed design decisions and help guide both the optimisation process and the specification of manufacturing tolerances.

Future direction: A systematic sensitivity analysis of the coil design parameters would provide insight into the relationship between geometry and electrical performance, identifying which dimensions most strongly influence inductance, resistance, and coupling. Such an analysis would also clarify whether the optimiser's solutions are robust or sit on narrow ridges of the design space, providing practical guidance for both the optimisation search bounds and manufacturing tolerance specifications.

4.8 Chapter Summary

This chapter has presented a complete MATLAB toolchain for rapid, air-core pre-design of inductive power transfer couplers. The toolchain comprises four sequential analysis tools: Tool 1 specifies optimal conductor dimensions for the target current and frequency; Tool 2 analyses individual coil electrical properties; Tool 3 evaluates coupling and efficiency over the user-defined misalignment envelope; and Tool 4 searches for optimal geometries using particle swarm optimisation.

The full toolchain was validated against Ansys Maxwell FEA over 864 evaluation points spanning 10 geometric test families, 4 Litz wire constructions, 4 operating frequencies, and 8 misalignment positions. The validation confirms error levels well-suited for design ranking: self-inductance is predicted with 0.1–1.3% typical error (3.1% maximum), mutual inductance with 0.5–5.5% typical error (9.8% maximum), and AC resistance with 0.7–5.8% typical error (12.8% maximum). These accuracy levels preserve design ranking, enabling the toolchain to confidently identify promising candidates for further FEA validation.

Despite the limitations discussed in Section 4.7, the toolchain achieves its strategic goal. The air-core simplifications enable rapid exploration of the design space and reliable design ranking, thereby focusing expensive FEA simulations only on genuinely promising candidates. The identified limitations do not invalidate the toolchain's usefulness; rather, they define its scope and

provide clear directions for future development, including analytical ferrite models, multi-receiver optimisation, Pareto front generation, extended geometry support, angular misalignment, and parameter sensitivity analysis.

With the analytical pre-filter established, the next chapter builds upon this foundation by integrating the toolchain into a complete five-stage design pipeline. Ferrite flux guides and aluminium shielding – deliberately excluded from the analytical stage – are introduced in Stage 3, where their influence on coupling, efficiency, and thermal behaviour is systematically evaluated through FEA. The optimised air-core geometries identified in this chapter thus serve as the starting point for the fully assembled, thermally validated coupler designs presented in Chapter 5. The toolchain's predictions are subsequently validated experimentally in Chapter 6, where the physical realisation of three of the five designs is described and compared against simulation results.

Chapter 5

Complete Methodology and Optimisation Results

This chapter presents the complete five-stage design pipeline and demonstrates its application to five magnetic coupler designs, designated *A-I*, *A-II*, *B*, *C*, and *D*. The pipeline integrates the analytical toolchain developed in Chapter 4 with finite-element validation, ferrite and shielding integration, electrical worst-case assessment, and thermal simulation. The results illustrate how the toolchain diagnoses proximity losses, guides Litz wire redesign, and enables systematic trade-offs between efficiency, cost, and weight across the full SAE J2954 Z1 misalignment envelope.

5.1 Five-Stage Design Pipeline: Overview

The methodology adopted in this work comprises five sequential stages, plus a concluding Design Closure step that validates the initial design choices. The pipeline is structured to balance computational speed with physical accuracy: fast analytical screening narrows the design space in Stage 1, while high-fidelity FEA simulations are reserved for the most geometrically sound candidates in Stages 2–5. Each stage introduces additional physical complexity, ensuring that computational effort is directed only where it is warranted.

The stages are as follows:

- **Stage 1: MATLAB Toolchain Pre-Filtering.** Using Tools 1–4 from Chapter 4, the designer specifies the Litz wire construction (Tool 1), analyses individual coil impedances (Tool 2), evaluates spatial coupling over the defined full SAE J2954 Z1 misalignment envelope (Tool 3), and performs global geometric optimisation (Tool 4). This stage functions as a high-speed geometric pre-filter: by operating in air-core mode, it rapidly identifies candidate geometries whose fundamental flux patterns are sound, before the computational and material complexity of ferrite and shielding are introduced. Stage 1 outputs optimised air-core geometries together with predicted total resistance, self-inductance, mutual inductance, coupling coefficient, and spatial efficiency statistics across the defined evaluation region.

- **Stage 2: FEA Air-Core Validation.** The geometric candidates from Stage 1 are re-simulated using Ansys Maxwell to verify the analytical predictions. The same spatial sweep is performed to extract mutual inductance and efficiency maps. Discrepancies are analysed to identify modelling errors (most notably proximity loss underestimation) and to guide toolchain refinement before ferrite material is introduced.
- **Stage 3: Ferrite and Shielding Integration.** Ferrite flux guides and aluminium back-shields are added to the FEA-validated air-core designs. A parametric sweep over a defined set of ferrite tile layouts is performed in Ansys Maxwell. The optimal layout for each coil size is selected based on a cost-performance metric evaluated over the full misalignment region, not merely at the aligned position. The chosen configuration is subsequently simulated across the complete SAE J2954 Z1 envelope.
- **Stage 4: Electrical Worst-Case Assessment.** With the complete magnetic assembly (coils + ferrite + shield), the full 3D misalignment sweep is re-evaluated to extract the minimum mutual inductance, $M_{\min}$. Using the series-series (SS) compensated system equations (Section 2.6.3), the RMS coil currents, peak voltages, and current densities are computed for this worst-case coupling condition. A saturation verification of the ferrite tiles under full power at $M_{\min}$ is performed to confirm linear operation. The sensitivity of system performance to resonance detuning is also discussed. The worst-case currents and peak voltages obtained serve to define the insulation requirements of the Litz wire. The worst-case current also serves to validate the initial current density estimate made in Stage 1.
- **Stage 5: Thermal Simulation.** The ohmic and core losses from the Maxwell worst-case simulation are mapped onto a 3D thermal model of the fully assembled pad. Natural convection and radiation boundary conditions are applied. The predicted maximum hotspot temperatures serve to define the thermal insulation requirements of the Litz wire, completing the electrical and thermal design of the conductor.
- **Design Closure (Validation and Iteration).** Stages 4 and 5 collectively serve as the validation checkpoints of the pipeline. The designer begins with an initial estimate of the maximum RMS current and an assumed bundle diameter or packing factor. These quantities affect proximity losses and self-inductance, which in turn influence the optimisation ranking. However, these are only guesses until Stages 4 and 5 provide quantified results. At that point, the designer can assess whether the computed maximum RMS current exceeds the initial estimate (requiring a redesigned Litz wire construction and a return to Stage 1), or whether the manufacturer-supplied bundle diameter differs significantly from the assumed value due to insulation thickness or packing factor. If this discrepancy is large enough to alter inductance or resistance substantially, the designer must restart from Stage 1 optimisation. In the present work, these validation checks are performed and discussed in Section 5.8, confirming that the initial assumptions were adequate and no iteration was required.

The following sections apply this pipeline systematically to the five candidate couplings.

5.2 Design Constraints and Specifications

The collaborative project within which this thesis is situated does not yet impose fixed physical constraints on the magnetic couplers (transmitter and receiver outer dimensions, air-gap range, etc.). Consequently, this chapter employs a self-consistent set of representative specifications to demonstrate and validate the design pipeline. The choices described below constitute a coherent experimental framework rather than fixed project requirements.

5.2.1 System-Wide Specifications

The firm system-level specifications adopted throughout this work are:

- **Power rating:** 3.7 kVA (SAE J2954 WPT1 class).
- **Operating frequency:** 85 kHz (nominal).
- **Compensation topology:** Series-series (SS). This topology is adopted as it is the one being implemented by the power electronics team and is already integrated into the toolchain (Chapter 4).
- **Control strategy:** The system is assumed to implement maximum-efficiency tracking, realised via a DC-DC converter, so that the ideal secondary load can be achieved at any operating position, as discussed in Chapter 2.

5.2.2 Misalignment Envelope and Evaluation Grid

In the absence of project-specific alignment constraints, the SAE J2954 Class Z1 misalignment envelope has been adopted as the evaluation standard: a vertical air gap of 100–150 mm, a lateral tolerance of ± 75 mm in the x -direction, and ± 100 mm in the y -direction. Angular tilt is not considered. By exploiting the geometric symmetry of the toolchain-implemented coils (as detailed in Chapter 4, Section 4.4.2), the full envelope is represented by its first quadrant, which is sampled on the grid defined in Table 5.1. This 60-position grid is used consistently in both the MATLAB optimisation (Stage 1) and the Ansys Maxwell validation sweeps (Stages 2–4).

Table 5.1: Evaluation grid for spatial sweeps (first quadrant of SAE J2954 Z1 envelope).

Axis	Min (mm)	Max (mm)	Step (mm)
x (lateral)	0	75	25
y (lateral)	0	100	25
z (air gap)	100	150	25

5.2.3 Remaining Design Inputs

The following design inputs are not fixed constraints but are treated as optimisation variables or engineering choices, and are discussed in detail in Section 5.3:

- **Coil geometries:** symmetric (500×500 mm, Tx and Rx) and asymmetric (Tx 500×500 mm, Rx 350×350 mm).
- **Litz wire constructions:** $300 \times \text{AWG32}$ and $1200 \times \text{AWG38}$, dimensioned using Tool 1.
- **Inductance ceiling:** $L < 200 \mu\text{H}$ (air-core). Large inductances require impractically small compensation capacitors, increase resonant voltage stresses, and render tuning highly sensitive to component tolerances. Therefore an upper bound was imposed on the air-core couplers so that the ferrite+shield couplers would also be limited. This is, however, an arbitrary initial guess, as no knowledge of how much the inductance will rise after ferrite integration is available during the air-core optimisation. If the final assembled inductances were found to drive the system into an unworkable capacitance or voltage regime, the ceiling would be tightened and Stage 1 re-run.

5.3 Stage 1: MATLAB Toolchain Pre-Filtering

Stage 1 constitutes the geometric pre-filter of the design pipeline. Operating without ferrite or shielding material, the MATLAB toolchain rapidly explores a large design space to identify coil geometries whose intrinsic flux patterns are physically sound. Only designs that emerge from this stage with fundamentally robust electromagnetic behaviour proceed to the computationally expensive FEA analyses of Stages 2–5. This section details the Litz wire dimensioning, the optimisation formulation, and the resulting five candidate designs.

5.3.1 Litz Wire Dimensioning (Tool 1)

The power electronics team provided an estimated peak RMS coil current of 30 A, derived from the 3.7 kVA power rating at a 400 V DC-link voltage, augmented by a safety margin to account for reactive circulation and potential overloads. A maximum current density of 3 A/mm^2 was adopted as the baseline conductor sizing constraint, placing this design in the conservative sub-range of natural convection cooling, as contextualised in Table 5.2.

Table 5.2: Typical current density recommendations for copper windings in IPT systems.

Cooling condition	Current density (A/mm^2)
Natural convection (conservative)	2–3
Natural convection (typical)	3–5
Forced air cooling	5–8
Liquid cooling	8–12

Tool 1 dimensioned two Litz wire constructions, both targeting an equivalent copper cross-section of approximately 9.6 mm^2 , so that their DC resistance per unit length is essentially identical. The strand diameter was selected as a fraction of the skin depth at 85 kHz ($\delta \approx 0.22 \text{ mm}$ for copper). The first construction adopts a safety factor $c_\delta = 1.0$, yielding 300 strands of AWG32. The second

adopts $c_\delta = 0.5$, yielding 1200 strands of AWG38. The rationale for this choice is deliberate: both constructions are adequate from a skin-effect standpoint, but proximity losses scale with the fourth power of the strand diameter. The coarse wire ($300 \times \text{AWG32}$) will therefore suffer significantly higher proximity losses in the dense magnetic field between adjacent turns, whereas the fine wire ($1200 \times \text{AWG38}$) largely suppresses this mechanism. By maintaining identical cross-sections, the comparison between these two constructions becomes a controlled experiment that isolates the impact of proximity effect. The complete output of Tool 1 is presented in Table 5.3.

Table 5.3: Litz wire construction parameters from Tool 1 for both the coarse-strand ($300 \times \text{AWG32}$) and fine-strand ($1200 \times \text{AWG38}$) candidates.

Parameter	$300 \times \text{AWG32}$	$1200 \times \text{AWG38}$
Safety factor c_δ	1.0	0.5
Strand diameter (mm)	0.202	0.101
Strand count	300	1200
Effective Cu area (mm^2)	9.61	9.62
Bundle outer diameter (mm)	4.6	4.6
DC resistance ($\text{m}\Omega/\text{m}$)	1.79	1.79
Skin factor F_{skin} (85 kHz)	0.0008	0.0002

5.3.2 Coupling Optimisation (Tool 4)

The optimisation objective is to maximise the spatially averaged coil-to-coil efficiency η_{max} (Eq. 2.24) over the 60-position evaluation grid of Table 5.1. Particle swarm optimisation (PSO) is employed with 20 particles and 50 iterations. The sole electrical constraint is the inductance ceiling $L < 200 \mu\text{H}$.

To ensure that the optimised geometries are not artifacts of a single stochastic run, the PSO was executed five times with different random seeds for each coupling variant. The best-performing run (by final objective value) was selected for each design. Across all five seeds, the final objective values differed by less than 0.05%, indicating that the optimisation consistently converged to the same region of the design space. The optimised parameters from the best run are reported in the following tables.

Five coupling variants are optimised as follows:

- **A-I:** Symmetric (500×500 mm), $300 \times \text{AWG32}$ wire. This design is a *non-ideal baseline*, produced by an earlier version of the toolchain that did not correctly account for proximity losses. The significance of this design is diagnostic rather than practical: it was produced by an optimiser that was blind to proximity-induced AC resistance, and therefore converged on touching turns (0 mm spacing) to maximise inductance and coupling. As demonstrated in Section 5.3.4, this winding geometry catastrophically inflates AC resistance. Following correction of the toolchain, A-I was re-evaluated with the final validated script to extract consistent comparative metrics.

- **A-II:** Symmetric (500×500 mm), $300 \times \text{AWG32}$ wire, re-optimised with the fully corrected toolchain. The pair (*A-I*, *A-II*) thus isolates the impact of toolchain fidelity: identical wire type and coil footprint, different loss model.
- **B:** Asymmetric (Tx 500×500 mm, Rx 350×350 mm), $300 \times \text{AWG32}$ wire, optimised with the corrected toolchain. The pair (*A-II*, *B*) isolates the effect of receiver size reduction under the same wire type.
- **C:** Symmetric (500×500 mm), $1200 \times \text{AWG38}$ wire, optimised with the corrected toolchain. Because proximity losses are dramatically reduced by the finer strands, the optimiser can exploit tighter turn spacing without a resistance penalty; the inductance ceiling becomes the sole binding constraint. The pair (*A-II*, *C*) isolates the effect of changing the Litz wire construction in a symmetric geometry, highlighting the geometric differences induced by the conductor choice.
- **D:** Asymmetric (Tx 500×500 mm, Rx 350×350 mm), $1200 \times \text{AWG38}$ wire, optimised with the corrected toolchain. The pair (*B*, *D*) shows the same wire-construction improvement in the asymmetric geometry.

5.3.3 Optimisation Results

Table 5.4 presents the optimised geometric parameters and predicted air-core electrical metrics for all five couplings. The following observations are of particular note:

- **A-I** uses touching turns (0 mm spacing) and 12 turns on both Tx and Rx. The total resistance per coil is $100.55 \text{ m}\Omega$, dominated by the AC proximity component ($R_{\text{ac}} = 66.28 \text{ m}\Omega$, versus $R_{\text{dc}} = 34.27 \text{ m}\Omega$). The average mutual inductance is $33.29 \text{ }\mu\text{H}$, yielding an average efficiency of 98.85%. This result represents the direct consequence of an optimisation that was blind to proximity losses: the geometry is electromagnetically aggressive but electrically penalised.
- **A-II**, produced by re-optimising the same wire type with the corrected toolchain, increases turn spacing to 4 mm and raises the turn count to 17. Despite the longer conductor (21.39 m per coil), total resistance falls to $69.38 \text{ m}\Omega$ because R_{ac} drops to $31.28 \text{ m}\Omega$. Self-inductance is virtually unchanged ($132.98 \text{ }\mu\text{H}$), mutual inductance increases to $38.75 \text{ }\mu\text{H}$ on average, and efficiency improves to 99.31%. This comparison with *A-I* constitutes a direct validation of the toolchain correction: accurate proximity modelling redirects the optimiser towards lower-loss geometries.
- **B** is the asymmetric coarse-wire design. The transmitter is geometrically identical to *A-II*; the receiver (350×350 mm, 11 turns, 4 mm spacing) has a conductor length of only 10.43 m and a resistance of $32.24 \text{ m}\Omega$. The reduced receiver aperture lowers mutual inductance to $M_{\text{avg}} = 16.98 \text{ }\mu\text{H}$ and reduces average efficiency to 98.93%.

- **C**, the symmetric fine-wire design, demonstrates the full potential of the 1200×AWG38 construction. With 20 turns and a 2 mm spacing, self-inductance reaches 199.78 μH – within 0.1% of the ceiling – while total resistance per coil is only 61.00 m Ω despite a conductor length of 25.59 m. Average mutual inductance is 57.43 μH , and average efficiency reaches 99.59%, the highest in the candidate set.
- **D**, the asymmetric fine-wire design, achieves $M_{\text{avg}} = 37.20 \mu\text{H}$ and an average efficiency of 99.39%, representing an improvement of approximately 0.46 percentage points over its coarse-wire counterpart **B**.

All five designs satisfy the air-core inductance constraint $L < 200 \mu\text{H}$. Three-dimensional coil visualisations generated by the MATLAB toolchain are shown in Fig. 5.1.

Table 5.4: Stage 1 optimised geometric parameters and air-core electrical metrics for all five coupler candidates.

Parameter	<i>A-I</i>	<i>A-II</i>	<i>B</i>		<i>C</i>	<i>D</i>	
	Tx/Rx	Tx/Rx	Tx	Rx	Tx/Rx	Tx	Rx
Geometric parameters							
Length (mm)	500	500	500	350	500	500	350
Width (mm)	500	500	500	350	500	500	350
Corner radius (mm)	150	180	180	105	190	190	105
Turn count	12	17	17	11	20	20	20
Turn spacing (mm)	0	4	4	4	2	2	0
Conductor length (m)	18.99	21.39	21.39	10.43	25.59	25.59	18.90
Electrical parameters (25 °C, 85 kHz)							
R_{dc} (m Ω)	34.27	38.10	38.10	18.50	45.78	45.78	33.92
R_{ac} (m Ω)	66.28	31.28	31.28	13.74	15.22	15.22	23.37
R (m Ω)	100.55	69.38	69.38	32.24	61.00	61.00	57.29
L (μH)	132.99	132.98	132.98	43.72	199.78	199.78	145.24
Q (-)	706.4	1023.7	1023.7	724.3	1749.3	1749.3	1354.1
Coupling metrics (25 °C, 85 kHz)							
M_{avg} (μH)	33.29	38.75	16.98		57.43	37.20	
k_{avg} (-)	0.2503	0.2914	0.2227		0.2874	0.2184	
η_{avg} (%)	98.85	99.31	98.93		99.59	99.39	

5.3.4 Proximity Loss Diagnosis

To quantify the impact of strand diameter on proximity-induced losses, the *A-I* geometry (touching turns, 12 turns, corner radius 150 mm) was selected as a deliberately challenging test case. Its dense winding creates a strong, non-uniform magnetic field at the conductor surfaces, making it the most proximity-loss-sensitive configuration in the candidate set. Both wire constructions were evaluated under identical geometric and excitation conditions using the toolchain.

Table 5.5 compares the resulting electrical parameters and proximity factors F_{prox} (as defined in Eq. 4.6). For the 300×AWG32 wire, $F_{\text{prox}} = 1.933$, indicating that proximity-induced losses are




Figure 5.1: Three-dimensional coil visualisations from the MATLAB toolchain: (a) *A-I*, (b) *A-II*, (c) *B*, (d) *C*, (e) *D*.

nearly twice the DC copper loss. For the $1200 \times \text{AWG38}$ wire, F_{prox} falls to 0.483 – a reduction of approximately a factor of four – reducing the total resistance from 100.55 mΩ to 50.83 mΩ and nearly doubling the quality factor from 706 to 1397 with no change in inductance.

Table 5.5: Proximity factor comparison for the *A-I* geometry evaluated with both Litz wire constructions.

Parameter	$300 \times \text{AWG32}$	$1200 \times \text{AWG38}$
R_{dc} (mΩ)	34.27	34.27
R_{ac} (mΩ)	66.28	16.56
R (mΩ)	100.55	50.83
L (μH)	132.99	132.99
Q (-)	706.4	1397.3
F_{skin}	0.0011	0.0001
F_{prox}	1.9331	0.4833

This result confirms that the finer Litz wire effectively suppresses proximity losses even under the most demanding winding conditions, justifying its use in couplings *C* and *D*, where the optimiser is subsequently free to reduce turn spacing without incurring a resistance penalty.

5.4 Stage 2: FEA Air-Core Validation

The five optimised air-core designs were re-simulated in Ansys Maxwell 2025 R2 using a 3D frequency-domain solver with a tetrahedral mesh. Both transmitter and receiver coils were excited at 30 A RMS, 85 kHz, and the same 60-position spatial sweep defined in Table 5.1 was applied. Two important clarifications apply to this excitation choice. First, the extraction of L , M , and R_{ac} is linear; these parameters are independent of excitation amplitude and therefore valid for all operating conditions, including asymmetric coupler configurations where the actual transmitter and receiver currents will differ. Second, the magnetic field magnitude visualisations presented in this section scale linearly with the 30 A excitation and are not representative of real operating flux densities under resonance; only the spatial distribution and relative patterns carry physical significance. Actual operating currents are determined in Stage 4.

Table 5.6 lists the spatially averaged air-core electrical metrics from Maxwell for all five couplings. The agreement with the MATLAB analytical predictions (Table 5.4) is very close, confirming that the toolchain's air-core models produce reliable results for these winding configurations. All five designs satisfy the electrical limits and proceed to the ferrite integration stage.

Table 5.6: Stage 2 air-core electrical metrics for all five couplings from Ansys Maxwell across the evaluation grid.

Parameter	<i>A-I</i>	<i>A-II</i>	<i>B</i>		<i>C</i>	<i>D</i>	
	Tx/Rx	Tx/Rx	Tx	Rx	Tx/Rx	Tx	Rx
Electrical parameters (25 °C, 85 kHz)							
R_{dc} (m Ω)	33.76	37.89	37.89	18.41	46.18	46.18	33.91
R_{ac} (m Ω)	64.19	32.24	32.24	14.03	15.14	15.14	20.55
R (m Ω)	97.95	70.13	70.13	32.44	61.32	61.32	54.46
L (μ H)	130.89	132.02	132.02	44.07	196.63	196.63	142.05
Q (-)	713.7	1005.6	1005.6	725.5	1712.7	1712.7	1638.9
Coupling metrics (25 °C, 85 kHz)							
M_{avg} (μ H)	30.44	35.63	15.57		52.67	33.97	
k_{avg} (-)	0.2326	0.2699	0.2040		0.2680	0.2032	
η_{avg} (%)	98.77	99.24	98.82		99.55	99.34	

Figure 5.2 presents the magnetic field distribution for two representative designs: the symmetric coupling *A-II* and the asymmetric coupling *B*. Cross-sectional side views (at the YZ -plane through the coil centres) are shown with a common colour scale to permit direct comparison. In the symmetric case (*A-II*), the flux pattern is mirror-symmetric about the mid-plane, with field lines bulging uniformly around the coil periphery. In the asymmetric case (*B*), the smaller 350×350 mm receiver confines the linked flux to a narrower aperture; the flux density immediately above the transmitter is consequently higher, but the net linked flux is reduced, consistent with the lower mutual inductance observed in Table 5.6.

Figure 5.3 shows the efficiency distribution over the first-quadrant evaluation grid for *A-II* and *B*, at air gaps of 100, 125, and 150 mm. The symmetric design (*A-II*) exhibits a broader and more




Figure 5.2: Magnetic field distribution (side view, YZ -plane) in the air-core configurations: (a) symmetric coupling $A-II$; (b) asymmetric coupling B .

uniform efficiency plateau, whereas coupling B shows a steeper gradient as the receiver offset increases, particularly at the maximum air gap. The spatial efficiency pattern closely mirrors that of the mutual inductance, as dictated by Eq. 2.24.




Figure 5.3: Coil-to-coil efficiency maps over the evaluation grid of the air-core configurations: (a) symmetric coupling $A-II$; (b) asymmetric coupling B .

5.5 Stage 3: Ferrite and Shielding Integration

Ferrite flux guides and aluminium back-shields are added to all five validated air-core designs. This stage serves a dual purpose: to quantify the coupling enhancement achievable with passive flux management, and to select a cost-effective material configuration through a spatially robust optimisation metric.

5.5.1 Design Standardisation Strategy

The ferrite layout introduces a large number of additional degrees of freedom (tile dimensions, coverage pattern, number of layers, and spacing from the coil plane) that, in principle, could be independently optimised for each coupling. To render the comparison between designs tractable and physically meaningful, the layouts were standardised according to the outer coil dimensions: a common set of layouts was defined for 500×500 mm coils (applicable to the transmitter of all couplings and the receiver of the symmetric designs) and a separate set for 350×350 mm coils (applicable to the receiver of the asymmetric designs B and D). Keeping the magnetic environment – ferrite coverage, tile type, thickness, and spacing – constant across couplings $A-II$,

B, *C*, and *D* isolates the contribution of different Litz wire constructions and winding geometries on electromagnetic performance. The aluminium shield dimensions (a 1×1 m plate on each side) were likewise kept constant.

5.5.2 Ferrite Tile Selection and Material Properties

Commercially available Mn-Zn ferrite tiles, specifically designed for inductive power transfer applications, were selected from an electronics supplier. The available tile geometries are 100×100 mm and 150×100 mm, with thicknesses of 5 to 20 mm (Fig. 5.4). The key magnetic properties of the selected material are summarised in Table 5.7. The initial relative permeability ($\mu_i = 3000 \pm 25\%$ at 23°C) remains approximately constant up to approximately 400 kHz (Fig. 5.5), well above the 85 kHz operating frequency; frequency-dependent permeability variation may therefore be neglected in this application.



Figure 5.4: Commercial Mn-Zn ferrite tiles geometries selected for coupler construction: (a) 100×100 mm, (b) 150×100 mm.

Table 5.7: Magnetic properties of the selected Mn-Zn ferrite material.

Property	Condition	Value
Initial permeability μ_i	23°C	$3000 \pm 25\%$
Core loss density P_{cv}	100 kHz, 200 mT, 100°C	330 kW/m^3
Curie temperature T_c	—	220°C
Saturation flux density B_{sat}	1200 A/m, 23°C	520 mT
	1200 A/m, 100°C	410 mT
Mass density ρ	—	4900 kg/m^3

5.5.3 Aluminium Shield Specification

The electromagnetic shield comprises a 1×1 m aluminium plate (alloy 1050, $\sigma \approx 37.7 \text{ MS/m}$) positioned on the far side of the ferrite on each pad. This footprint exceeds the maximum coil outer dimension (500×500 mm) and conforms to common practice in SAE J2954 reference designs. The skin depth in this alloy at 85 kHz is approximately 0.28 mm; a 1 mm thick plate therefore attenuates the field to negligible levels within a few skin depths, providing effective shielding




Figure 5.5: Relative permeability versus frequency for the selected Mn-Zn ferrite (typical curve from supplier datasheet).

without excessive eddy-current loss. Shield thickness and standoff distance from the ferrite were determined by the parametric study described in the following section.

5.5.4 Standardised Ferrite and Shield Parameters

A comprehensive per-parameter optimisation of ferrite and shield geometry for each of the five couplings would be computationally prohibitive. Instead, a parametric FEA study was conducted on a single representative assembly – coupling *A-II* with a *Type IV* ferrite layout – at the centred position ($z = 150$ mm, $x = y = 0$ mm). The following parameters were swept independently: ferrite thickness (5, 10, 15, 20 mm), ferrite standoff from the coil (0–25 mm in 5 mm steps), shield thickness (0.5, 1.0, 1.5, 2.0 mm), and shield standoff from the ferrite (0–25 mm in 5 mm steps). Total resistance, self-inductance, mutual inductance, and efficiency were extracted for each combination. The values selected from this analysis – summarised in Table 5.8 – were then applied uniformly across all five couplings.

Table 5.8: Standardised ferrite tile and aluminium shield assembly parameters applied uniformly across all five coupler designs.

Ferrite tile	Value (mm)	Aluminium shield	Value (mm)
Thickness	5	Thickness	1
Standoff (coil-to-ferrite)	15	Standoff (ferrite-to-shield)	20

The 5 mm ferrite thickness provides effective flux guidance without excessive weight or material cost. The 15 mm coil-to-ferrite gap reduces eddy-current losses in the tile faces while still

enhancing the coupling coefficient. A 1 mm shield is sufficient to contain stray fields at 85 kHz, and the 20 mm ferrite-to-shield standoff limits the inductance-reducing counter-flux effect of the shield. Maintaining these values constant across all couplings ensures that observed performance differences are attributable solely to the winding geometry and Litz wire construction, not to variations in the magnetic environment.

5.5.5 Coupling *A-I* as a Non-Ideal Baseline

coupling *A-I* is deliberately treated as a non-ideal baseline and is not included in the standardised layout comparison. A reduced-ferrite arrangement, designated *Type X* (Fig. 5.6), was intentionally applied to this coupling to observe the performance limits of a suboptimal design when magnetic flux guidance is minimised. This choice reflects the design intent: coupling *A-I* already suffers from elevated copper losses due to its touching-turn winding geometry; the application of a sparse ferrite layout amplifies this deficit and permits the combined consequences of poor winding design and insufficient flux concentration to be quantified and contrasted with the optimised designs.




Figure 5.6: Ferrite tile layout for the coupler of coupling *A-I* (*Type X*).

5.5.6 Per-Coupling Ferrite Layout Definition

500×500 mm coils. Eight tile layouts (*Type I–VIII*) were defined for the 500×500 mm coil footprint. *Type I–V* consist of full-coverage plates built from combinations of 150×100 mm and 100×100 mm tiles, with some variants incorporating a central aperture. *Type VI–VIII* cover only the active winding area (Fig. 5.7).

350×350 mm receiver coils. A separate set of six layouts (Fig. 5.8) was defined for the 350×350 mm receiver of couplings *B* and *D*. The transmitter of both asymmetric couplings (*B* and *D*) inherits the layout selected for the 500×500 mm coils.




Figure 5.7: Ferrite tile layouts for 500×500 mm coils: (a) *Type I*, (b) *Type II*, (c) *Type III*, (d) *Type IV*, (e) *Type V*, (f) *Type VI*, (g) *Type VII*, (h) *Type VIII*.




Figure 5.8: Ferrite tile layouts for 350×350 mm receiver coils: (a) *Type I*, (b) *Type II*, (c) *Type III*, (d) *Type IV*, (e) *Type V*, (f) *Type VI*.

5.5.7 Spatial Layout Selection Methodology and Cost-Performance Metric

All ferrite layouts were simulated in Ansys Maxwell across the full first-quadrant SAE J2954 Z1 evaluation grid (Table 5.1). Layout selection was therefore based on performance averaged over the entire misalignment region, not on the centred position alone. This approach ensures that the chosen layout performs robustly across the real operational envelope rather than being optimised for a single, favourable alignment.

For each layout, the spatial averages of mutual inductance M_{avg} and maximum efficiency η_{avg} were extracted from the 60-position sweep. The ferrite material cost was computed from the tile count and sizes using the reference prices in Table 5.9. The layout selected for each coupling maximises the efficiency increment per unit ferrite cost (Eq. 5.1):

$$\frac{\Delta\eta_{\text{avg}}}{\mathcal{C}} = \frac{\eta_{\text{avg, ferrite+shield}} - \eta_{\text{avg, air-core}}}{\text{ferrite cost}}, \quad (5.1)$$

where $\mathcal{C}$ is a representative cost proxy based on supplier pricing at the time of design, expressed in euros per tile. Because absolute material prices are subject to market fluctuation and supplier-specific agreements, $\mathcal{C}$ is best interpreted as a normalised measure of material volume rather than a precise economic figure. The $\Delta\eta_{\text{avg}}$ values obtained in this way are spatially robust averages that are representative of real-world system operation across the full Z1 envelope.

Table 5.9: Reference unit costs for ferrite tiles and aluminium shield material used in the layout cost-performance metric.

Component	Dimensions	Unit cost
Ferrite tile	100×100×5 mm	13.72€/unit
Ferrite tile	150×100×5 mm	24.21€/unit
Aluminium plate	1×1×0.001 m	3.32€/kg

The complete results for all layouts are compiled in Appendix B, which lists, for every layout candidate and every coupling, the total resistance, self-inductance, quality factor, spatially averaged mutual inductance, average coupling coefficient, maximum efficiency across the evaluation region, total ferrite cost, and the resulting cost-performance metric. For the asymmetric couplings *B* and *D*, only the receiver-side layout was varied in the parametric study; the transmitter layout was held fixed at the *Type IV* configuration selected for the 500×500 mm symmetric coils. This choice ensures that any performance difference between the asymmetric and symmetric designs is attributable exclusively to the smaller receiver aperture and its tile arrangement, without introducing a simultaneous change on the transmitter side.

Table 5.10 lists the selected layout for each coupling, the associated spatial efficiency and mutual inductance increments, and the cost-performance ratio.

The negative $\Delta\eta_{\text{avg}}$ for coupling *A-I* warrants explicit interpretation. The special *Type X* layout provides only sparse flux guidance; consequently, the eddy-current losses induced in the 1 mm aluminium shield and the partial flux cancellation introduced by the shield in the vicinity of the

Table 5.10: Selected ferrite layouts and cost-performance metrics for each coupling, spatially averaged over the evaluation grid.

coupling	Tx layout	Rx layout	ΔM_{avg} (%)	$\Delta \eta_{\text{avg}}$ (p.p.)	Total ferrite (cost €)	$\Delta \eta / \mathcal{C}$ ($\times 10^6$ % per €)
<i>A-I</i>	<i>Type X</i>	<i>Type X</i>	4.14	−0.48	581.04	−826.0
<i>A-II</i>	<i>Type IV</i>	<i>Type IV</i>	112.74	0.10	1078.16	92.8
<i>B</i>	<i>Type IV</i>	<i>Type IV</i>	77.52	0.07	787.64	88.9
<i>C</i>	<i>Type IV</i>	<i>Type IV</i>	111.09	0.14	1078.16	129.9
<i>D</i>	<i>Type IV</i>	<i>Type IV</i>	76.54	0.10	787.64	127.0

coil exceed the marginal coupling benefit of the reduced ferrite coverage. The net effect is a measurable reduction in average coil-to-coil efficiency relative to the air-core baseline. This result is physically meaningful and confirms that a suboptimal winding geometry (touching turns, coarse wire) combined with an insufficient ferrite layout can actively degrade system performance rather than improving it. This ferrite arrangement was retained for coupling *A-I* precisely because it makes this degradation mechanism observable and quantifiable.

5.5.8 Effect of Ferrite and Shield on Electrical Parameters

Following layout selection, the complete assembly (coils + ferrite + aluminium shield) was simulated across the full 60-position evaluation grid. The reported AC resistance R_{ac} for the fully assembled couplers reflects a true pad-to-pad loss budget, comprising three distinct contributions:

1. Litz wire copper losses (skin and proximity effects in the winding),
2. Ferrite core losses (hysteresis and eddy-current losses in the tiles),
3. Aluminium shield eddy-current losses (induced in the passive shield plate).

The efficiency values reported in Table 5.11 therefore represent pad-to-pad efficiency, not coil-only efficiency. This distinction is critical when comparing the assembled results with the air-core values of Table 5.6.

The addition of ferrite increases both mutual and self-inductance; the aluminium shield slightly reduces both due to its counter-flux effect. The net outcome is a moderate improvement in spatially averaged efficiency for couplings *A-II*, *B*, *C*, and *D*: the ferrite consistently raises the average coupling coefficient by approximately 3–8 percentage points and increases the efficiency by 0.1–0.2 absolute percentage points. The asymmetric couplings *B* and *D* benefit proportionally more from the ferrite, as the flux concentration effect of the tile array compensates partially for the reduced aperture of the 350×350 mm receiver.

Figure 5.9 shows the assembled magnetic field distribution (side view) for couplings *A-II* and *B*. Compared with the air-core counterparts in Fig. 5.2, the ferrite visibly concentrates flux towards the receiver, increasing linked flux, while the shield attenuates the field on the back surface. This redirection effect is particularly pronounced for coupling *B*, where the ferrite partially compensates for flux that would otherwise escape laterally beyond the smaller receiver aperture.

Table 5.11: Stage 3 electrical metrics with ferrite and aluminium shield.

Parameter	<i>A-I</i>	<i>A-II</i>	<i>B</i>		<i>C</i>	<i>D</i>	
	Tx/Rx	Tx/Rx	Tx	Rx	Tx/Rx	Tx	Rx
R_{dc} (m Ω)	33.76	37.89	37.89	18.41	46.18	46.18	33.91
R_{ac} (m Ω)	104.47	89.90	89.90	30.86	43.82	43.82	52.18
R (m Ω)	138.23	127.79	127.79	49.27	90.00	90.00	80.04
L (μ H)	158.76	217.05	217.05	65.06	323.26	323.26	211.70
Q (-)	613.4	908.0	908.0	705.4	1918.1	1918.1	1412.7
M_{avg} (μ H)	31.70	75.80	27.64		111.18	59.97	
k_{avg} (-)	0.1996	0.3465	0.2393		0.3425	0.2361	
η_{avg} (%)	98.29	99.34	98.89		99.68	99.44	

5.6 Stage 4: Electrical Worst-Case Assessment

With the fully assembled magnetic structure defined, Stage 4 determines the electrical operating point under the most demanding coupling condition within the SAE J2954 Z1 envelope and verifies that the ferrite material remains in its linear operating regime.

The complete magnetic assembly (coils + ferrite + shield) was simulated across the full 60-position evaluation grid to obtain the spatial distribution of mutual inductance M . In a series-series compensated system operating under resonance, the minimum mutual inductance M_{min} defines the worst-case design point, because it forces the highest coil currents and peak voltages for a given output power. This worst-case sensitivity to M_{min} is a direct consequence of the SS topology's load-independence: the resonant frequency is set by the pad inductances and their compensation capacitors alone, and cannot adapt to changes in coupling. While this property guarantees resonance stability across the Z1 envelope (Section 5.2), it also means that the system must be dimensioned to deliver full rated power even at the minimum coupling point without saturating the ferrite or exceeding conductor insulation ratings. Within the SAE J2954 Z1 envelope, M_{min} consistently occurs at the maximum lateral offset and maximum air gap, i.e., at the corner position ($x = 75$ mm, $y = 100$ mm, $z = 150$ mm).




Figure 5.9: Magnetic field distribution (side view) with *Type IV* ferrite and aluminium shield: (a) coupling *A-II*; (b) coupling *B*.

5.6.1 Worst-Case Electrical Analysis

At the worst-case position, the coupling is assumed to deliver the full rated output power $P_{\text{out}} = 3.7$ kW to the load. Perfect resonance is maintained, and the optimal load resistance $R_{L,\text{opt}}$ from Eq. 2.23 (Chapter 2) is applied. The RMS coil currents follow from the power balance and reflected impedance relations:

$$I_2 = \sqrt{\frac{P_{\text{out}}}{R_{L,\text{opt}}}}, \quad I_1 = \frac{(R_2 + R_{L,\text{opt}}) I_2}{\omega M_{\text{min}}}. \quad (5.2)$$

For asymmetric couplings (B and D), the different coil self-inductances and resistances produce unequal transmitter and receiver currents. The resonant compensation capacitances are computed from the fully assembled self-inductances (Table 5.11) using Eq. 2.25. The peak coil and capacitor voltages are equal at resonance (Eq. 5.3):

$$\hat{V} = \sqrt{2} \omega L I, \quad (5.3)$$

and these values determine the voltage rating of the compensation capacitors and the insulation class of the Litz wire.

Table 5.12 presents the computed worst-case electrical parameters for all five couplings. The following observations merit discussion:

- **Current density validation:** For both Litz wire constructions, the maximum RMS coil current remains below the initial 30 A estimate, confirming that the current density does not exceed the 3 A/mm² design limit. The initial conductor sizing is therefore adequate.
- **Asymmetric current distribution:** For symmetric couplings $A-I$, $A-II$, and C , transmitter and receiver currents are approximately equal. For coupling B , the smaller receiver coil carries a significantly higher current (27.13 A) than the transmitter (17.02 A), a consequence of the impedance mismatch between the two pad sizes. This imbalance concentrates dissipation on the secondary side and is verified in the thermal simulations of Stage 5. Coupling D exhibits a smaller but qualitatively similar imbalance.
- **Peak voltages and insulation requirements:** Peak coil and capacitor voltages reach 2790 V for coupling B and up to 3439 V for coupling D . These values assume ideal resonance at the worst misalignment and constitute the baseline specification for compensation capacitor voltage ratings and Litz wire insulation class. In practice, detuning and transient effects may generate higher instantaneous voltages; the Stage 4 values are therefore used as minimum design benchmarks. Because the required capacitance values (as low as 10.85 nF for coupling C) are not readily available as single commercial components, series-parallel capacitor networks will be required to achieve both the target capacitance and adequate voltage margin.

Table 5.12: Stage 4 worst-case electrical parameters for all five couplings under rated output of 3.7 kW.

Parameter	<i>A-I</i>		<i>A-II</i>		<i>B</i>		<i>C</i>		<i>D</i>	
	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx
$M_{\min}$ (μH)	18.06		42.40		14.75		62.41		32.15	
$R_{L,\text{opt}}$ (Ω)	9.65		22.65		4.89		33.33		16.09	
C (nF)	22.08	22.08	16.15	16.15	16.15	53.89	10.85	10.85	10.85	16.87
I_{rms} (A)	19.60	19.32	12.68	12.61	17.02	27.13	10.42	10.39	14.08	14.96
Peak $\hat{V}_L, \hat{V}_C$ (V)	2350	2316	2079	2067	2790	1333	2544	2537	3439	2347

5.6.2 Ferrite Saturation Verification

A critical verification step is the confirmation that the Mn-Zn ferrite tiles operate within the linear regime under full-power conditions. Unlike the electrical worst-case, which occurs at the minimum-coupling corner position ($x = 75, y = 100, z = 150$ mm), the magnetic worst case for the ferrite tiles occurs at the aligned, minimum air-gap position ($x = 0, y = 0, z = 100$ mm). At this position the flux density threading the tiles is at its maximum: the coils are fully overlapping and the air gap is smallest, so the ferrite carries the highest flux for a given excitation current. Symmetrically, this aligned position also corresponds to $M_{\max}$, the maximum mutual inductance across the entire Z1 evaluation grid. Following the same procedure as Stage 4, $M_{\max}$ was used in conjunction with the SS system equations (Eq. 5.2) to compute the transmitter and receiver RMS currents at this position under rated output power of 3.7 kW. These position-specific currents were then applied as excitation amplitudes in the Maxwell simulation, ensuring that the extracted $B_{\max}$ reflects the true worst-case flux density in the ferrite under realistic operating conditions rather than a conservative overestimate. The relevant material limit is the saturation flux density at elevated temperature: $B_{\text{sat}} = 410$ mT at 100°C (Table 5.7). Exceeding this threshold would invalidate all inductance predictions derived from linear FEA simulations and would cause a sharp increase in core losses.

The results are presented in Table 5.13. For all five designs, the simulated peak flux density $B_{\max}$ remains well below the 410 mT saturation limit, with the highest observed value of 105.29 mT occurring on the receiver side of coupling *B*. The 5 mm tile thickness and the standardised 15 mm coil-to-ferrite standoff are therefore confirmed to be adequate for full-rated operation across the entire SAE J2954 Z1 envelope. All couplings operate within the linear magnetic regime and no ferrite saturation occurs under the design conditions.

The $B_{\max}$ values reveal several important physical distinctions. The highest flux density is found in coupling *B*, on the receiver side (105.29 mT), which is significantly higher than its transmitter (35.65 mT) and also exceeds all other designs. This is a direct consequence of the smaller 350×350 mm receiver aperture: the flux from the 500×500 mm transmitter is concentrated into a reduced ferrite area, increasing the local flux density. The receiver of coupling *D*, also a 350×350 mm coil, shows a similar but lower elevation (70.40 mT), reflecting its lower mutual inductance and current requirements.

coupling *A-I*, which employs the sparse *Type X* layout, exhibits moderate flux densities of 75.21 mT on the transmitter and 69.64 mT on the receiver. These values, while higher than the full-coverage symmetric designs (*A-II* and *C*), are still well below the saturation limit. The sparse ferrite coverage does concentrate flux through a reduced cross-section, but the effect is less pronounced than initially anticipated. The highest values are driven more by the aperture reduction in the asymmetric receiver than by the tile coverage pattern.

Table 5.13: Peak ferrite flux density $B_{\max}$ (mT) at the aligned minimum air-gap position for all five couplings.

Parameter	<i>A-I</i>		<i>A-II</i>		<i>B</i>		<i>C</i>		<i>D</i>	
	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx
Ferrite layout	<i>Type X</i>		<i>Type IV</i>		<i>Type IV</i>		<i>Type IV</i>		<i>Type IV</i>	
$B_{\max}$ (mT)	75.21	69.64	36.74	33.11	35.652	105.29	40.23	38.79	47.66	70.40
B_{sat} (mT)	410 (at 100 °C)									

5.6.3 Resonance Detuning Sensitivity

The foregoing analysis assumes perfect resonance, i.e., the operating frequency satisfies Eq. 2.20 exactly. In practice, component tolerances (typically $\pm 5\%$ for capacitors, $\pm 10\%$ for inductors) and temperature drift perturb the resonant condition. Taking coupling *C* as a representative example, the nominal fully assembled inductances are $L_1 = L_2 = 323.26 \mu\text{H}$ (Table 5.11), requiring $C_1 = C_2 \approx 10.85 \text{ nF}$ for resonance at 85 kHz.

Figure 5.10 illustrates the frequency response of a series *LC* circuit in the vicinity of resonance. The resonant point is characterised by a minimum in impedance magnitude and a zero phase angle crossing; it is inherently sensitive to component variations, since a small shift in either *L* or *C* moves the operating point off-resonance, introduces reactive impedance, and increases both current and apparent power. For example, for coupling *C*, a $+5\%$ inductance deviation shifts resonance to approximately 82.8 kHz; operation at the fixed 85 kHz would result in an inductive load, an input VA increase of more than 20%, and an efficiency penalty of 1–2 percentage points. A corresponding -5% deviation shifts resonance to 87.2 kHz with capacitive consequences of similar magnitude.

Furthermore, because the transmitter and receiver tanks are independent circuits, their component tolerances may shift resonance frequencies asymmetrically. This asymmetric detuning induces additional reactive power circulation between the two sides, placing further stress on the power semiconductor switches. A fixed-frequency operating strategy is therefore not robust against typical component variations. A phase-locked loop frequency tracking mechanism, or an adaptive tuning scheme based on switchable capacitor banks, is required to maintain near-resonant operation. The design of such control is within the scope of the parallel power electronics effort. The worst-case electrical stresses determined in Section 5.6.1 provide the boundary conditions for that controller: at the corner misalignment position of the Z1 envelope, the expected coil currents and peak voltages are fully characterised.




Figure 5.10: Typical frequency response of a series LC circuit near resonance (magnitude in red, phase in blue).

5.7 Stage 5: Thermal Simulation

The ohmic and core losses computed in the Maxwell worst-case scenario of Stage 4 were transferred to a 3D thermal model in Ansys Icepak. The model incorporates the complete pad geometry – Litz wire bundle, ferrite tile array, aluminium shield plate, and surrounding air volume – with loss density distributed uniformly over each conductor volume. For a fully transposed Litz wire bundle, this uniform distribution is a conservative approximation that slightly overestimates peak hotspot temperatures. The thermal boundary conditions follow the procedure of the Ansys Icepak tutorial for natural convection and radiation [4]; forced convection and direct thermal management are not assumed.

Temperature distributions for coupling *A-II* (Figs. 5.11) and coupling *B* (Fig. 5.12) are presented separately for the coil, ferrite, and aluminium shield sub-domains. Colour scales differ between couplings; within each figure, the three sub-domain plots share a common scale to permit component-level comparison.




Figure 5.11: Steady-state temperature distribution for coupling *A-II* under worst-case losses: (a) Litz wire coils, (b) ferrite tiles, (c) aluminium shield.




Figure 5.12: Steady-state temperature distribution for coupling *B* under worst-case losses: (a) Litz wire coils, (b) ferrite tiles, (c) aluminium shield.

The maximum hotspot temperatures for all components and all five couplings are summarised in Table 5.14. Coupling *A-I* exhibits the highest temperatures across all components, with a coil hotspot of 52.5 °C, ferrite at 42.8 °C, and shield at 39.4 °C—a direct consequence of the elevated copper losses associated with its touching-turn winding and coarse Litz wire. For the asymmetric coupling *B*, the secondary-side coil reaches 47.5 °C (versus 42.9 °C for the transmitter), consistent with the higher receiver RMS current identified in Stage 4. In all cases, temperatures remain well below the Curie temperature of the ferrite (220 °C) and within the thermal ratings of the Litz wire insulation, as confirmed in Section 5.8.

Table 5.14: Stage 5 maximum component hotspot temperatures (°C) for all five couplings under worst-case power dissipation with natural convection cooling.

Component	<i>A-I</i>		<i>A-II</i>		<i>B</i>		<i>C</i>		<i>D</i>	
	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx
Coil	51.3	52.5	38.22	39.0	42.9	47.5	31.9	32.2	35.1	38.6
Ferrite	41.1	42.8	33.5	35.1	36.1	41.2	30.4	31.3	33.3	34.2
Shield	39.4	39.4	32.8	32.8	34.4	34.3	29.8	29.8	30.7	31.6

Based on these results, the 300×AWG32 Litz wire (couplings *A-I*, *A-II*, *B*) must be rated for a maximum operating temperature of at least 52.5 °C, while the 1200×AWG38 construction (couplings *C*, *D*) must tolerate a minimum of 38.6 °C. These values are initial estimates derived under the natural convection assumption. The final mechanical enclosure, including any supports or ventilation constraints, may reduce passive cooling effectiveness and may lead to higher temperatures at each component, especially at nominal output power.

5.8 Design Closure

The five-stage pipeline was initiated with a set of engineering assumptions: the Litz wire constructions, the maximum RMS current, the current density limit, and the inductance ceiling. The quantitative outputs of Stages 4 and 5 now permit a systematic retrospective validation of each assumption.

5.8.1 Litz Wire Electrical and Thermal Validation

Table 5.15 compares the minimum wire requirements derived from the pipeline with the corresponding rated values from the selected Litz wire suppliers.

Table 5.15: Validation of Litz wire electrical and thermal requirements against supplier specifications.

Parameter	300×AWG32	1200×AWG38
Required		
Insulation voltage (V)	≥ 2790	≥ 3439
Temperature rating (°C)	≥ 52.5	≥ 38.6
Bundle diameter (mm)	4.6	4.6
Supplier specifications		
Insulation material	Double-layer silk	Single-layer polyimide
Breakdown voltage (V)	3000	6000
Temperature rating (°C)	155.0	180.0
Bundle diameter (mm)	4.8	4.8
Associated cost (€ per m)	3.202	3.234

Both constructions satisfy the electrical and thermal requirements with substantial margins. The 1200×AWG38 wire is particularly well-suited to future higher-power applications, given its 6000 V breakdown voltage. The estimated bundle diameter of 4.6 mm deviates from the supplier's measured 4.8 mm by less than 5%; this discrepancy has a negligible effect on inductance and resistance predictions, and does not warrant repeating the Stage 2 FEA simulations. Were the deviation larger Stage 2 onward would need to be revisited.

Figure 5.13 shows both Litz wire constructions. The coarse-wire couplings *A-I*, *A-II*, and *B* were physically assembled and experimentally characterised within the scope of this thesis; the results are presented in Chapter 6. The fine-wire couplings *C* and *D* remain as validated simulation designs available for future prototype work.

5.8.2 Inductance Limit Validation

The air-core inductance ceiling of $L < 200 \mu\text{H}$ was imposed at Stage 1 as a design closure safeguard to prevent the ferrite-augmented inductances from reaching values at which compensation circuit practicality is compromised. As detailed in Section 5.2.3, exceeding this headroom would push the required compensation capacitances into a sub-10 nF regime dominated by parasitic reactances, while simultaneously elevating resonant peak voltages towards or beyond the insulation limits of the Litz wire.

Table 5.11 confirms that the most inductance-sensitive design, coupling *C*, reaches a fully assembled self-inductance of approximately $323 \mu\text{H}$ – a factor of 1.62 above the air-core value of $200 \mu\text{H}$. The corresponding compensation capacitance is approximately 10.85 nF, which, while achievable with commercial high-voltage film capacitors, requires a series-parallel network to




Figure 5.13: Photograph of the two purchased Litz wire constructions.

simultaneously satisfy the capacitance and voltage rating specifications. Had the air-core inductance ceiling been relaxed, this value would fall below 10 nF – the practical threshold at which resonance tuning becomes critically sensitive to PCB parasitic inductances and to manufacturing tolerances in both the capacitor bank and the winding geometry. The 200 μ H limit thus proves to be a well-motivated constraint that was identified as necessary from first principles and is here retrospectively confirmed by the Stage 3 inductance values.

5.8.3 Final Coupler Cost Summary

Table 5.16 presents the bill-of-materials cost breakdown for each coupler design, covering the three physical components: Litz wire, ferrite tiles, and the aluminium shield. These figures are based on the reference unit prices listed in Table 5.9 and should be treated as indicative estimates rather than fixed procurement costs, since ferrite tile prices in particular are subject to significant market fluctuation depending on supplier, order volume, and availability at the time of purchase.

Across all five designs, the ferrite tiles represent by far the dominant cost contributor, accounting for 39–43% of the total for the symmetric designs and 72–79% for the finer-wire variants where the coil cost is comparatively modest. The aluminium shield, shared across all configurations at a fixed 1×1 m plate per coupler, contributes a negligible 22.60€ per coupler. The Litz wire cost scales directly with conductor length and is consequently higher for the fine-wire designs (*C*, *D*) despite their lower per-metre price, owing to the greater turn counts demanded by the inductance ceiling. The asymmetric designs (*B*, *D*) benefit from a substantially cheaper receiver coil and receiver ferrite array, reducing total system cost by approximately 25% relative to their symmetric counterparts while accepting the coupling penalty analysed in preceding sections.

5.9 Critical Assessment and Future Directions

The five-stage pipeline presented in this chapter demonstrates the practical utility of the complete design methodology through the systematic design and optimisation of five couplers. The

Table 5.16: Bill-of-materials cost breakdown (€) for each coupler design.

Component	<i>A-I</i>	<i>A-II</i>	<i>B</i>		<i>C</i>	<i>D</i>	
	Tx/Rx	Tx/Rx	Tx	Rx	Tx/Rx	Tx	Rx
Coil	60.80	68.49	68.49	33.40	82.75	82.75	61.11
Ferrite	290.52	539.08	539.08	248.56	539.08	539.08	248.56
Shield				45.20			
Total cost	747.85	1260.34	934.72		1288.85	976.70	

results establish a coherent set of designs whose electromagnetic, electrical, and thermal behaviour is fully characterised across the SAE J2954 Z1 misalignment envelope. However, the pipeline has several limitations that stem from choices made to keep the five-coupling comparison tractable. These limitations affect the generalisability of the results, the scalability of the methodology, and the depth of the validation. The following subsections synthesise these limitations and identify concrete directions for future development.

5.9.1 Standardised Ferrite Layouts

The most significant limitation of the pipeline concerns the standardisation of ferrite layouts by coil dimension. As discussed in Section 5.5.1, a common set of layouts was defined for 500×500 mm coils and a separate set for 350×350 mm coils. This standardisation ensured that observed performance differences were attributable to winding geometry and wire type, not to variations in the magnetic environment. However, this pragmatic choice likely left marginal efficiency gains unexplored relative to per-coupling custom optimisation.

The magnitude of these potential gains remains unknown. A per-coupling optimisation could, in principle, tailor the ferrite coverage pattern, tile thickness, and spacing to the specific flux distribution of each winding geometry, potentially improving coupling and efficiency beyond the values reported in Table 5.11. The standardisation strategy was a necessary compromise to keep the comparison tractable—optimising ferrite layouts independently for each of the five couplings would have required significantly increased simulation effort—but it does mean that the designs presented here are not necessarily optimal in the absolute sense. They are optimal within the constrained set of layouts considered.

Future direction: Extending the ferrite layout selection to per-coupling custom optimisation through automated FEA sweeps or surrogate-assisted optimisation would quantify the true cost of standardisation and capture any marginal efficiency gains left unexplored in this work. This would require a more sophisticated optimisation framework that can handle the additional degrees of freedom introduced by ferrite tile placement, but the cost-performance metric $\Delta\eta/\mathcal{C}$ introduced in Section 5.5.7 provides a natural foundation for such an extension.

5.9.2 Automation and Scalability Limitations

The current pipeline relies on manual setup of Ansys Maxwell simulations for each candidate design. While the 60-position spatial sweep is handled through a parametric sweep in x , y , and z , the overall process across Stages 2–4 still requires significant manual intervention. The coils, ferrite layouts, and shield geometries must be defined and positioned manually; the meshing strategy must be configured for each combination; and the results—resistance, inductance, coupling, losses, and field distributions—must be extracted and organised for further analysis. This manual process is time-consuming and error-prone, and it limits the speed and scalability of the methodology.

The implication is that while the pipeline is effective for a small number of candidate designs (five in this work), it would become impractical for larger design spaces or more extensive parametric studies. The manual steps in geometry definition, simulation setup, and data extraction are bottlenecks that prevent the methodology from being applied to, for example, a full Pareto front optimisation or a comprehensive sensitivity analysis. The worst-case electrical and thermal analyses suffer from the same limitations, as they require additional simulations with position-specific currents and loss mappings that must be manually configured and extracted.

Future direction: Automating the remaining manual steps through Python or MATLAB scripting—including coil and ferrite geometry generation, simulation setup, and results extraction for the 60-position sweeps, as well as the handoffs between Maxwell and Icepak—would dramatically reduce turnaround time for Stages 2–4. This would enable more extensive parametric studies, more thorough validation, and the application of the methodology to larger design spaces. The scripting infrastructure developed for this work provides a foundation for such automation, but further development is required to make it fully robust and user-friendly.

5.9.3 Thermal Simulation Assumptions

Stage 5 assumed natural convection and radiation only, with the pads exposed to open air. This assumption was appropriate for the initial thermal validation—it provides a conservative estimate of temperatures under the simplest cooling scenario. However, in a practical application, the pads will be enclosed in plastic housings with limited ventilation. This enclosure effect could reduce passive cooling effectiveness, potentially increasing component temperatures beyond the values predicted in Table 5.14.

The thermal margins observed (maximum coil temperature of 52.5 °C, well below the 155 °C insulation rating) suggest the designs are robust to some reduction in cooling effectiveness. However, the magnitude of the temperature increase due to enclosure is unknown. Without experimental validation under representative conditions, the thermal safety margin cannot be quantified with confidence. This is a gap between the simulation results and the real-world installation conditions that must be addressed before the designs can be deployed in production.

Future direction: Experimentally validating thermal behaviour within a representative enclosure, including forced convection or active cooling strategies, would bridge the gap between open-air

simulations and real-world installation conditions. This would provide the data needed to quantify the true thermal safety margin and, if necessary, to refine the thermal design of the pads.

5.9.4 SS Topology Dependency

The efficiency expressions underpinning Stage 4 are strictly valid for the series-series compensation topology adopted throughout this work. The choice of SS was motivated by the parallel power electronics effort—which was already developing an SS-compensated system.

This dependency limits the generalisability of the electrical analysis. The coupler designs presented in this chapter are optimised for SS compensation, meaning that their performance under other compensation topologies may be suboptimal. The electrical worst-case assessment, the current and voltage calculations, and the efficiency estimates are all specific to SS.

Future direction: Extending the electrical analysis to support series-parallel, parallel-series, parallel-parallel, and LCC topologies would broaden the applicability of the methodology and enable consistent topology comparison. This would allow designers to select the optimal compensation topology for their specific application and to compare coupler designs on a consistent basis across topologies.

5.9.5 Absence of Tolerance Analysis

The pipeline does not include a systematic tolerance analysis to assess how variations in key parameters affect performance. While the resonance detuning sensitivity analysis in Section 5.6.3 touches on this—examining the impact of $\pm 5\%$ inductance deviations on resonance—this is only a partial treatment. A comprehensive analysis would also need to consider the tolerances of the compensation capacitors, which are typically $\pm 5\%$ to $\pm 10\%$ for commercial high-voltage film capacitors. Variations in capacitance directly shift the resonant frequency, introduce reactive power circulation, and increase losses.

Furthermore, the efficiency expressions derived in Section 2.6.3 assume ideal load conditions, i.e., that the equivalent load resistance seen by the secondary side is exactly $R_{L,\text{opt}}$. In practice, the DC-DC converter implementing MEPT may not achieve the ideal load perfectly due to control inaccuracies, dynamic transients, or limited resolution. If the load deviates from the optimal value, the efficiency drops, and the system performance no longer follows the ideal curves presented in this chapter.

The absence of such analysis means that the designs presented here are optimised for nominal conditions and ideal control. Their sensitivity to component tolerances and control non-idealities is unknown. This is a significant gap for production-ready systems, where component variations and control imperfections are inevitable.

Future direction: A systematic sensitivity analysis incorporating inductance and capacitor tolerances, as well as deviations from the ideal load condition, would provide insight into the robustness of the designs under realistic operating conditions. This would enable designers to assess

the expected performance spread and to specify appropriate component tolerances and control accuracy requirements.

5.9.6 Absence of SAE J2954 Reference Geometry Comparison

No quantitative comparison with the SAE J2954 reference pad geometries has been performed within this chapter. The primary objective of this work was the development and demonstration of the five-stage methodology itself, not a benchmarking study against the standard's reference designs. The self-consistent framework adopted here—comparing the five optimised couplings against each other to isolate the effects of wire type and geometry—was a deliberate choice to keep the study tractable and to focus on the internal validity of the pipeline.

Consequently, the performance claim that the methodology "may outperform the reference designs" (Section 3.7) remains unverified. While the optimised couplings *C* and *D* show high simulated efficiencies within the chosen design space, it is not known whether they would exceed the performance of an SAE J2954 reference pad under identical conditions. This does not undermine the methodology's utility, but it does mean that the results should be interpreted as internally consistent design examples rather than as a demonstrated improvement over the state of the art.

Future direction: A systematic comparison between the optimised designs and the SAE J2954 reference pad geometries would provide a valuable external validation of the methodology. Such a study would answer the question of whether the designs genuinely improve upon an established reference or merely perform well within the constraints chosen for this work.

5.10 Chapter Summary

This chapter has presented and applied the five-stage design pipeline to five magnetic coupler candidates, taking each from an initial analytical geometry through FEA air-core validation, ferrite and shielding integration, electrical worst-case assessment, and thermal simulation, before closing the loop with a retrospective validation of all initial design assumptions. The results establish a coherent, internally consistent set of designs whose electromagnetic, electrical, and thermal behaviour is fully characterised across the SAE J2954 Z1 misalignment envelope.

The comparison between coupling *A-I* and coupling *A-II* demonstrated the necessity of accurate loss modelling, with the corrected toolchain increasing turn spacing from 0 to 4 mm and improving average efficiency from 98.85% to 99.31%. The standardised ferrite layout comparison showed that spatially averaged metrics, rather than centred-position optimisation, ensure reliable performance under realistic parking conditions. The *Type IV* layout was consistently optimal across couplings *A-II*, *B*, *C*, and *D*, with efficiency increments of 0.13–0.15 absolute percentage points per unit ferrite cost.

Electrical worst-case assessment confirmed that all five designs remain within the linear ferrite regime, with peak flux densities well below the 410 mT saturation limit. Thermal simulations predicted maximum hotspot temperatures well within the limits of both the Litz wire insulation and the ferrite Curie temperature, with the highest temperatures occurring for coupling *A-I* (coil hotspot

52.5 °C). The design closure step validated the initial assumptions on current density, inductance ceiling, and conductor insulation requirements.

Despite the limitations discussed in Section 5.9 the pipeline achieves its primary objective: to provide a systematic, standards-compliant pathway for the design of high-performance magnetic couplers. The comparison between coupling *A-I* and *A-II* validates the necessity of accurate loss modelling; the spatially robust ferrite optimisation demonstrates the value of averaging over the full misalignment envelope; and the design-closure framework confirms the importance of retrospective validation of initial assumptions. The methodology is fit for its intended purpose, and the identified limitations provide clear directions for future refinement.

Of the five designs, couplings *A-I*, *A-II*, and *B*—all wound with the 300×AWG32 Litz wire—are ready for physical realisation. The two fine-wire designs (*C* and *D*) represent a higher-performance trajectory that is fully validated at the simulation level and available for future prototype development. Chapter 6 takes the three coarse-wire couplings from design to hardware, describing their fabrication, experimental characterisation, and direct comparison against the predictions established here.

Chapter 6

Practical Implementation and Validation

6.1 Introduction

The design methodology presented in the preceding chapters – comprising the analytical toolchain (Chapter 4) and the five-stage pipeline (Chapter 5) – relies heavily on simulations and theoretical models. While these tools have been validated against finite-element analysis, their ultimate credibility depends on comparison with physical measurements. This chapter bridges that gap by describing the practical implementation and experimental validation of three of the five coupler designs.

From the five optimised couplings (*A-I*, *A-II*, *B*, *C*, *D*), three were selected for physical construction: *A-I*, *A-II*, and *B*. All three share the same Litz wire construction (300×AWG32), which was readily available. The remaining two designs (*C* and *D*) use a finer 1200×AWG38 wire that could not be sourced within the project timeline – their validation remains a topic for future work.

The experimental campaign included the design and fabrication of a non-magnetic test jig, the winding of coils using custom 3D-printed holders, and the preparation of Litz wire terminations by soldering. Measurements were performed using a four-wire Kelvin method for DC resistance, a Bode 100 vector network analyser for individual coil characterisation, and the same instrument for mutual inductance measurements over the SAE J2954 misalignment envelope.

6.2 Mechanical Test Jig Design

Magnetic field measurements are easily perturbed by nearby metal parts, which can introduce eddy currents or flux distortion. It is therefore standard practice to build test fixtures for inductive power transfer couplers from non-magnetic materials. Following this principle, the entire jig was constructed using a wooden base, acrylic sheets for structural frames and sliding surfaces, and 3D-printed plastic materials for all other components needed.

6.2.1 Overall Structure and Three-Axis Movement

The jig rests on a wooden base, which provides a flat, stable and electrically inert platform. Four vertical 3D-printed threaded rods are fixed to the base at the corners, supporting the entire upper assembly. By turning the 3D-printed nuts on them, the receiver sandwich is raised or lowered, adjusting the air gap z . The exact gap is measured with a digital calliper between the facing surfaces of the transmitter and receiver sandwiches, and the actual distance between coil surfaces is estimated from that measurement.

Horizontal movement is achieved by a free-sliding mechanism. The receiver sandwich sits on acrylic rails that allow sliding in the x direction. The transmitter sandwich, sliding on perpendicular rails, provides movement in the y direction. By pushing the sandwiches by hand, any combination of x and y misalignment within the SAE J2954 envelope (and even beyond) can be set. Millimetre scales on the rails give an approximate reading, and the exact position is verified with a digital calliper before each measurement. This simple design avoids backlash and ensures good repeatability. The aluminium shields are placed on top of the top acrylic sandwich and on the bottom of the bottom one, as discussed in Chapter 5.5.

Figure 6.1 shows a 3D Computer-Aided Design (CAD) representation of the complete jig with the coordinate axes represented as well as the transmitter and receiver sandwiches. The top aluminium sheet (from the receiver side) is absent to better show the components.




Figure 6.1: 3D CAD model of the test jig. The arrows indicate the x , y and z directions of movement.

6.2.2 Tx and Rx Sandwich Assembly

Both the transmitter and the receiver are built as a sandwich of three acrylic sheets, each with a distinct function. The top sheet (5 mm thick) holds the coil. Instead of gluing the winding directly, custom 3D-printed corner holders, matching the exact corner shape of each coil geometry, are screwed into this sheet. The Litz wire is wound around these holders (to be presented in Section 6.3). The middle sheet (10 mm thick) provides structural rigidity and carries the mechanical load. The bottom sheet (5 mm thick) serves as a carrier for the ferrite tiles. The three sheets are clamped together with 3D-printed screws, forming a single rigid unit that can be lifted and handled easily.

Figure 6.2 shows an exploded view of the sandwich for the transmitter of coupling *A-II*, with the three acrylic layers, 3D-printed corner holders, and ferrite tiles visible alongside other fixation components.




Figure 6.2: Exploded view of a sandwich (transmitter of coupling *A-II*).

Ferrite tiles are relatively expensive, so the same set of tiles is reused across all couplings. Different ferrite layouts are obtained simply by rearranging the tiles on the bottom sheet. The tiles are temporarily fixed with double-sided adhesive tape to the middle sheet, allowing rapid reconfiguration.

6.2.3 Thermal Considerations

All 3D-printed components (threaded rods, nuts, corner holders) are produced by Fused Deposition Modelling (FDM) 3D-printing using standard Polylactic Acid (PLA) filament. PLA is easy to print, sufficiently stiff for mechanical loads, and completely non-magnetic. However, its glass transition temperature T_g lies in the range 55–65°C – above this temperature it softens and can deform under sustained stress.

In the low-power validation measurements reported here (mainly signal-level excitation), the coils remain essentially at room temperature, so PLA is perfectly safe. Nevertheless, for coupling *A-I* the thermal simulation in Chapter 5 predicted a coil temperature of approximately 52.5°C. When enclosed in a finished pad with limited airflow, the temperature could rise further. This value is already close to the lower end of the PLA T_g range, and any increase in power or ambient temperature could cause deformation.

Therefore, while PLA is adequate for the low-power validation presented in this chapter, for future high-power tests the coil holders should be reprinted in a filament with higher thermal resistance, such as Polyethylene Terephthalate Glycol (PETG) ($T_g \approx 80^\circ\text{C}$). The other 3D-printed parts (threaded rods, nuts, alignment frames) experience negligible thermal stress and can remain in PLA. For the present experiments, all results were obtained with PLA holders at ambient temperature, and no deformation was observed.

6.3 Coil Holders and Modular Winding

The test jig was designed with maximum modularity in mind. The middle and bottom acrylic sheets of each sandwich are identical for all couplers – only the top acrylic sheet, which carries the coil, needs to be exchanged. In order to ease the winding process of coil designs, interchangeable 3D-printed corner holders that define the turn spacing, corner radius, and overall shape were developed. By changing the top sheet and the corresponding holders, the same sandwich can be reconfigured from coupling *A-I* to *A-II* or to the receiver of coupling *B* in a matter of minutes.

A total of five coils were wound and built for this validation: two identical coils for coupling *A-I* (Tx and Rx), two identical coils for coupling *A-II* (Tx and Rx), and one additional coil for the receiver of coupling *B* (the transmitter of *B* is the same as the *A-II* coil). Thus, only three distinct corner holder designs were required, corresponding to the three geometries presented in Figure 6.3.

The holder for the *A-I* coils (leftmost design) was a simple one-piece guide with no separation between turns. After winding the first coil, this design proved unsatisfactory: if too much tension was applied, the wire tended to force out of shape; if too little tension was used, the unsupported straight sections sagged and spread apart. Therefore a redesign was implemented for the *A-II* and *B* coils, consisting of two pieces that clamp together with plastic screws. The lower piece has grooves that precisely match the wire diameter, ensuring that each turn is held without the need for tensioning. The upper piece applies a uniform pressure across the coil, preventing any movement or sagging during winding and handling. This design proved to be more effective in maintaining the intended geometry of the coil throughout the manufacturing process.

Even with this redesign, some sagging was still noticeable on the unsupported straight sections. To address this, small additional holders were placed along these sections to maintain the correct wire spacing. Furthermore, a more complex structure was added to manage the connecting wires coming out of the coil, keeping them organised and preventing mechanical stress on the terminations. Figures 6.4a and 6.4b illustrate the process. The first shows the two-piece holder for the *A-II* transmitter with the coil almost fully wound and the wire maintains its shape without




Figure 6.3: CAD models of the three corner holders: (a) *A-I*, (b) *A-II*, (c) *B-Rx*.

external tensioning. The second shows the same coil after all turns have been seated and mounted to the acrylic sheet with all the additional holders.




Figure 6.4: Winding steps for the *A-II* transmitter: (a) final stage of the winding process, (b) completed coil.

Because the grooves precisely match the wire diameter, winding becomes extremely straightforward. The Litz wire is unwound directly from the supply spool and pressed into the grooves of the lower piece turn by turn, starting from the outermost groove and working inward. Once all turns are seated, the upper piece is placed over the wire and secured, and the entire assembly is screwed onto the top acrylic sheet of the sandwich. An extra length of wire (approximately 75 cm) was left on each side of the coil to provide flexibility for future connections and sufficient margin for soldering and crimping.

6.4 Soldering and Crimping of Litz Wire

All five coils use the same Litz wire construction (300×AWG32). The insulation of this wire is a double-layered silk, which behaves more like a fabric than a common plastic insulation. When stripped, the silk tends to continue unraveling along the wire. To prevent this, a short piece of Polyimide tape (PI tape) was wrapped around the wire at the desired stripping boundary, stopping any further separation. Approximately 15 mm of wire was exposed to match the chosen crimp connector.

A solder pot was filled with lead-free solder and set to approximately 400°C. At this temperature the solder slightly damages the insulation close to the stripped end but the PI tape protects the region beyond. The exposed strands were fluffed (spread apart) and flux was applied. The end was then dipped vertically into the solder pot for a few seconds. After removal and cooling, the termination was cut in order to examine the cross-section bonding. Three typical outcomes are illustrated in Figure 6.5.

- **Bad solder (Figure 6.5a):** Little to no flux was used, and the temperature was too low. The wire appeared soldered, but after cooling and applying pressure, an outer shell of solder broke away, revealing shiny, untouched copper.
- **Almost perfect (Figure 6.5b):** A tiny central area still showed bare copper, meaning the solder had not fully penetrated the bundle.
- **Perfect solder (Figure 6.5c):** The cross section revealed no weak spots, covering every strand uniformly.




Figure 6.5: Comparison of soldered terminations: (a) poor, (b) almost perfect, (c) perfect.

After soldering, any residual flux was carefully scraped away from the termination area to avoid poor electrical contact. A manual crimper was then used to attach an eye-shaped crimp terminal to each prepared wire end. Figure 6.6 shows one completed coil termination with the eye crimp and heat-shrink tubing in place. This process was repeated for all five coils.

Proper termination is critical. If the soldering or crimping is defective, some strands may carry no current at all because they remain electrically isolated from the rest of the bundle. Even a



Figure 6.6: Completed termination of a coil.

well-made solder joint and crimp can still become a localised hot spot, but the effect is greatly reduced compared to a poor termination. The true test of the crimping quality will be reflected in the DC and AC resistance measurements presented in the following sections.

6.5 Assembled Coils and Test Jig

After winding, soldering and crimping, each coil was mounted onto its dedicated top acrylic sheet. Figure 6.7 shows one representative top sheet for each of the three distinct geometries.






Figure 6.7: Top acrylic sheets with wound coils: (a) A-I, (b) A-II, (c) B-Rx.

Once the top sheet was ready, it was stacked together with the middle acrylic sheet, the bottom acrylic sheet, and the ferrite layout to form a complete sandwich. Figure 6.8a shows a fully assembled sandwich for the transmitter of coupling A-II. Finally, a transmitter and receiver sandwich were installed on the respective sliding rails. Figure 6.8b presents a photograph of the complete test jig in the laboratory, with the receiver sandwich placed at a representative position.

This view complements the 3D CAD model shown earlier (Figure 6.1) and demonstrates the actual implementation of the three-axis movement system.



(a)



(b)

Figure 6.8: Physical implementations of: (a) assembled sandwich (transmitter of coupling A-II) and (b) test jig in the laboratory (receiver side aluminium shield removed for visibility).

6.6 Measurement of DC Resistance

DC resistance is a fundamental electrical parameter of a coil. However, for the coils considered here, the DC resistance is expected to be extremely low (tens of milliohms). Even with a sophisticated ohmmeter, measuring such small resistances accurately is challenging since the voltage drop across the test leads and contact resistances can be comparable to, or even larger than, the voltage drop across the coil itself.

6.6.1 Two-Wire versus Four-Wire Method

In a standard two-wire ohmmeter (Figure 6.9a), a test current is forced through the same pair of leads that are used to measure the voltage drop. Any resistance in the leads or at the contacts (R_{wire}) adds directly to the measurement, making the reading less accurate for milliohm-level measurements.

The four-wire Kelvin method (Figure 6.9b) overcomes this problem by separating the current-forcing and voltage-sensing paths. Two leads carry the test current I_{test} from the current source to the coil. Two separate sense leads are connected directly across the coil terminals to measure the voltage drop V_{sense} . Because the sense leads carry negligible current, the measured voltage excludes the voltage drops across the force leads and their contacts. The DC resistance is then simply $R_{\text{dc}} = V_{\text{sense}}/I_{\text{test}}$.




Figure 6.9: Comparison of (a) two-wire and (b) four-wire Kelvin measurement methods. Adapted from [2].

6.6.2 Measurement Results

All DC resistance measurements were performed using a Tektronix DMM4050 digital multi-meter in its four-wire ohms mode. The measurements were repeated three times per coil, and the average value is reported. The ambient temperature was monitored throughout and remained within $25 \pm 2^\circ\text{C}$.

Table 6.1 presents the measured DC resistance for each coil alongside the predicted values from Ansys Maxwell stated in Chapter 5.

Table 6.1: Measured DC resistance compared with simulation results.

Source	<i>A-I</i>		<i>A-II</i>		<i>B</i>	
	Tx	Rx	Tx	Rx	Tx	Rx
Maxwell	33.76	33.76	37.89	37.89	37.89	18.41
Measured	38.8	38.5	42.5	43.4	42.5	22.4
δR_{dc} (%)	14.9	14.0	12.2	14.5	12.2	21.7

6.6.3 Discussion of Discrepancies

The measured DC resistance is higher than the simulated values, reaching over 20% error for the smaller coupler. The main identified factors that may contribute to this difference are listed below:

- **Extra conductor length:** As noted in Section 6.3, an additional 75 cm of wire was left on each side of every coil. This extra length yields approximately 2.6–2.7 m Ω for each coil. For the receiver of coupling *B*, where the total error in Table 6.1 is approximately 4.0 m Ω , this accounts for roughly 65–70% of the discrepancy. The remaining error must be attributed to the next two factors.
- **Contact resistance of soldered terminations:** The crimped eye terminals and solder joints, even when well executed, introduce additional contact resistance. For high-strand-count

Litz wire, micro-gaps or incomplete solder wetting in a few strands can produce measurable resistance. A “perfect” solder (Figure 6.5c) reduces but does not eliminate this effect.

- **Temperature variation:** Temperature deviations from 25°C affect copper resistivity. A $\pm 2^\circ\text{C}$ variation introduces a resistance shift of approximately 2–3%, which for coupling B Rx amounts to 0.6–1.3 m Ω . This is non-negligible given the tight measurement margins.

The combined effect of these three sources reasonably explains the observed 12–22% errors in Table 6.1. The overall agreement validates that the conductor length calculations and termination procedures are sound for design purposes.

6.7 Measurements of Individual Coils

Before performing coupling measurements, each coil was characterised in isolation to extract its self-inductance L and total resistance R at the operating frequency. In order to validate the Ansys Maxwell simulations, both the air-core version and the ferrite+shield version of each coil were characterised.

6.7.1 Measurement Approach

All measurements were performed with a Bode 100 vector network analyser from Omicron Lab. The instrument can be configured as a one-port impedance analyser, measuring magnitude $|Z|$ and phase θ over a user-defined frequency range. It also includes an automatic model-fitting function: the user selects an equivalent circuit from a library of eight passive models (combinations of resistors, inductors and capacitors), and the device estimates the component values by minimising the RMS relative error (RMSRE) between the measured impedance response and the model’s transfer function.

The objective was to characterise each coil at the intended operating frequency of 85 kHz. Near this frequency, the coils can be accurately described as a frequency-dependent resistor in series with a constant inductor: $Z(f) = R(f) + j2\pi fL$. The inductance L remains essentially constant for air-core coils and also for the ferrite+aluminium assembly, since the ferrite tiles do not saturate at the excitation level used. The resistance $R(f)$, however, varies with frequency due to skin and proximity effects, as discussed in previous chapters.

A wide frequency sweep (e.g., from 70 kHz to 90 kHz) would introduce a systematic error: the auto-fit would attempt to represent the frequency-dependent $R(f)$ by an ideal constant resistor, and part of the resistance variation could be absorbed into the estimated inductance. Therefore, a narrow sweep from 84.5 kHz to 85.5 kHz with a total of 400 sample points was chosen. Over this 1 kHz bandwidth, the resistance variation is negligible, so the simple series $R+L$ model is appropriate.

The Bode 100 was operated in one-port impedance mode, measuring both magnitude and phase. The built-in auto-fit was set to **Model B**, presented in Figure 6.10. This model accounts for the inevitable parasitic capacitance of the coil. The device then outputs the fitted values of R , L and the parallel capacitance.




Figure 6.10: Impedance measurement and model fitting for a representative coil. (a) Measured magnitude (red) and phase (blue). (b) Model B and the estimated parameters.

Before each measurement series, the Bode 100 was calibrated using its built-in open/short/load routine. This removes the effects of the Bayonet Neill-Concelman (BNC) cables and the test fixture up to the calibration plane, which is located at the end of the BNC connectors. However, to physically connect to the coil terminals, short alligator clips were attached beyond this plane, as illustrated in Figure 6.11. The parasitic resistance introduced by these clips (mainly due to the small contact area and additional wire length) is not compensated by the instrument's calibration. To estimate this parasitic contribution, a manual zeroing step was performed: the two alligator clips were shorted together and the same narrow-sweep impedance measurement was executed. The estimated resistance was recorded and subsequently subtracted as a bias term from all coil measurements. Repeated zeroing measurements showed contact resistance variability of ± 2.5 mΩ due to pressure-dependent contact area and oxidation, introducing a similar uncertainty into each individual coil measurement. The following tables present the compensated values, with this contact uncertainty acknowledged as a secondary contribution to the total measurement error.

The coils were placed on the test jig far from any metallic objects to avoid external electromagnetic interference. The results obtained with this procedure are reported in the following sections.

6.7.2 Measurements of Air-Core Coils

Table 6.2 presents the measured total resistance and self-inductance for each air-core coil, together with the predictions from Ansys Maxwell. The measured metrics are presented as X_{meas} whereas the simulated metrics are presented as X_{sim} . The relative error (δ) is calculated as in Equation 4.9.

The measured inductances agree closely with the simulations (within 0.3% for the 500 mm coils and 3.2% for the smaller B_{Rx} coil), confirming that the winding process and the 3D-printed holders successfully reproduce the designed geometries. The measured total resistances show larger




Figure 6.11: Measurement setup showing the reference planes (before and after calibration and after the manual zeroing step).

deviations, particularly for the *A-II* and *B* designs. A detailed discussion of the possible sources of errors is deferred to Section 6.9.

Table 6.2: Measured total resistance and self-inductance for air-core coils.

	<i>A-I</i>		<i>A-II</i>		<i>B</i>	
	Tx	Rx	Tx	Rx	Tx	Rx
R_{sim} (m Ω)	97.95	97.95	70.13	70.13	70.13	32.44
L_{sim} (μ H)	132.99	132.99	132.98	132.98	132.98	43.72
R_{meas} (m Ω)	108.40	107.13	88.55	85.54	88.55	49.27
L_{meas} (μ H)	132.99	133.33	132.57	131.22	132.57	45.12
δR (%)	10.67	9.37	26.27	21.98	26.27	51.88
δL (%)	0.00	0.26	-0.31	-1.32	-0.31	3.20

6.7.3 Measurements with Ferrite and Aluminium Shield

After the air-core characterisation, the same coils were assembled into sandwiches with the respective ferrite layouts and aluminium shielding. Table 6.3 presents the measured total resistance and self-inductance for each coupler following the same naming convention as Table 6.2.

The measured resistance is significantly higher than in the air-core case, especially for the smaller receiver, where it reaches 40% relative error. The inductance also presents deviations, particularly for the ferrite layouts used in the *A-II* and *B* couplers, which employ full ferrite coverage and are therefore more sensitive to small variations in the coil geometry and the exact positioning of the ferrite tiles. The discussion of the possible sources of errors is deferred to Section 6.9.

Table 6.3: Measured total resistance and self-inductance for fully assembled coils.

	<i>A-I</i>		<i>A-II</i>		<i>B</i>	
	Tx	Rx	Tx	Rx	Tx	Rx
R_{sim} (m Ω)	138.23	138.23	127.79	127.79	127.79	49.27
L_{sim} (μ H)	158.76	158.76	217.05	217.05	217.05	65.06
R_{meas} (m Ω)	154.27	169.86	152.48	154.40	154.40	69.01
L_{meas} (μ H)	159.80	157.51	193.31	192.17	192.17	60.38
δR (%)	11.60	22.88	19.32	20.82	20.82	40.06
δL (%)	0.66	-0.79	-10.94	-11.46	-11.46	-7.19

6.8 Full Coupling Measurements

With the individual coils characterised, the next step was to measure the mutual inductance M of the three couplings (*A-I*, *A-II*, *B*) at various misalignment positions. The objective was to validate the Ansys Maxwell models and the mechanical test jig by comparing measured values with simulated predictions. Any discrepancies help identify inaccuracies and sources of error in the modelling or experimental setup.

6.8.1 Measurement Approach

The mutual inductance estimation follows the series-aiding and series-opposing method, as described in Section 2.4.3 of Chapter 2.

The Bode 100 was configured in one-port impedance mode with the same narrow sweep used for the individual coil characterisation. For each misalignment position, the two sandwiches (with ferrite and aluminium shield) were placed in the test jig at the required coordinates. The series-aiding/series-opposing connections were made manually by swapping the terminals of one coil and each measurement was repeated three times and the average value recorded.

Ideally, the full 60-point grid defined in Table 5.1 should be measured, but achieving a step size of 25 mm with high repeatability on the physical jig is time-consuming and prone to small positioning errors. Therefore, only the eight extreme positions were evaluated. This still captures the worst-case coupling conditions and provides a stringent test of the model.

6.8.2 Measured Mutual Inductance

Table 6.4 presents the measured mutual inductance M_{meas} , the Maxwell prediction M_{sim} of the coupling, and the relative error (δ) for each of the eight positions.

The measured mutual inductance agrees well with the Maxwell predictions for coupling *A-I* (errors below 4%). For couplings *A-II* and *B*, the errors are larger (up to 10%), which will be discussed in Section 6.9. Nevertheless, the overall trend is captured correctly, confirming that the test jig, the assembly process, and the measurement method are reliable. The discrepancies will be further analysed in the following chapters to refine both the models and the experimental procedure.

Table 6.4: Measured mutual inductance compared with simulation results.

Position (mm)	<i>A-I</i>			<i>A-II</i>			<i>B</i>		
	$M_{\text{meas}} (\mu\text{H})$	$M_{\text{sim}} (\mu\text{H})$	$\delta M (\%)$	$M_{\text{meas}} (\mu\text{H})$	$M_{\text{sim}} (\mu\text{H})$	$\delta M (\%)$	$M_{\text{meas}} (\mu\text{H})$	$M_{\text{sim}} (\mu\text{H})$	$\delta M (\%)$
(0,0,100)	50.07	49.95	-0.23	103.11	114.92	10.28	38.23	42.61	10.27
(75,0,100)	41.01	40.69	-0.77	89.97	98.07	8.27	32.93	36.16	8.92
(0,100,100)	36.18	35.46	-2.03	81.62	86.52	5.66	31.60	31.60	8.66
(75,100,100)	29.65	28.77	-3.05	68.54	72.35	5.26	24.87	26.20	5.09
(0,0,150)	31.13	30.16	-3.23	64.21	70.44	8.84	23.45	26.01	9.84
(75,0,150)	27.20	26.17	-3.94	55.76	61.56	9.42	20.27	22.43	9.64
(0,100,150)	24.03	23.51	-2.18	51.08	55.28	7.60	18.60	19.88	6.43
(75,100,150)	20.75	20.18	-2.85	45.20	47.71	5.25	15.64	16.89	7.39

6.8.3 Extrapolated Coupling Coefficient and Maximum Efficiency

Using the measured mutual inductance M_{meas} from Table 6.4 and the self-inductances L_1, L_2 and total resistances R_1, R_2 from the ferrite+shield measurements (Table 6.3), the coupling coefficient k and the maximum achievable efficiency η_{max} (for ideal resonant compensation and optimal load) were computed for each of the eight extreme positions. Table 6.5 presents the relative error of k and the absolute error of η_{max} for each position and each coupling.

Table 6.5: Relative and absolute errors of the coupling coefficient and maximum efficiency for each test coupling.

Position (mm)	<i>A-I</i>		<i>A-II</i>		<i>B</i>	
	$\delta k (\%)$	$\Delta \eta (\text{p.p.})$	$\delta k (\%)$	$\Delta \eta (\text{p.p.})$	$\delta k (\%)$	$\Delta \eta (\text{p.p.})$
(0,0,100)	-1.37	0.166	-4.84	0.127	3.75	0.314
(75,0,100)	-0.74	0.191	-6.87	0.123	1.43	0.336
(0,100,100)	-1.52	0.199	-9.75	0.111	0.72	0.371
(75,100,100)	-2.24	0.216	-9.97	0.121	-3.73	0.376
(0,0,150)	-3.54	0.241	0.28	0.254	5.28	0.553
(75,0,150)	-4.19	0.262	0.91	0.291	4.76	0.628
(0,100,150)	-2.31	0.331	-1.11	0.296	1.17	0.630
(75,100,150)	-2.81	0.361	-3.70	0.305	1.93	0.757

The errors are small for coupling *A-I*, confirming that the test jig and the measurement procedure are reliable. For couplings *A-II* and *B*, the errors are larger (up to around -10% for k and 0.76 percentage points for η_{max}), which is consistent with the mutual inductance discrepancies discussed in the previous section.

6.9 Critical Assessment and Future Directions

The experimental validation presented in this chapter confirmed the fundamental design ranking established in simulation, but also revealed several important limitations that affect the accuracy of the measured quantities and the generalisability of the results. These limitations can be grouped into

three categories: measurement-induced errors (phase sensitivity and contact resistance variability), ferrite-related uncertainties (permeability tolerance and positioning), and mechanical and thermal limitations of the test jig (slack in the 3D-printed structure and PLA thermal constraints). Beyond these, two additional scope limitations were identified: the low-power nature of the validation and the unimplemented fine-wire designs (couplings *C* and *D*). The following subsections synthesise these limitations and identify concrete directions for future development.

6.9.1 Measurement-Induced Errors

The one-port impedance measurement with the Bode 100 exhibited significant phase sensitivity, particularly for the small receiver of coupling *B*, where the measured resistance error reached 52%. This phase sensitivity arises because the impedance phase angle is extremely close to 90° at 85 kHz – a direct consequence of the high quality factors of the air-core coils. A tiny absolute error in phase measurement translates into a large relative error in the extracted total resistance. This effect is compounded by contact resistance variability from the alligator clips, which introduces additional uncertainty on the order of 2–5 m Ω .

These measurement uncertainties propagate to the computed maximum efficiency. For coupling *B*, the efficiency error of up to 0.76 percentage points is comparable to the efficiency difference between competing designs, meaning that the measurement uncertainty could obscure real differences in performance.

The implications for the present work are twofold. First, the ranking consistency observed across all measurements – coupling *A-II* consistently outperforming *A-I* – is robust despite these errors, confirming that the relative ordering of designs is preserved. Second, the absolute values of resistance and efficiency should be interpreted with caution, particularly for the smaller coils where the errors are most pronounced. Future work should adopt a measurement method that is inherently less sensitive to phase errors.

Future direction: A forced-resonance measurement setup using a series *LC* circuit with a low-ESR capacitor should be built to characterise high-*Q* couplers. By measuring the –3 dB bandwidth of the impedance magnitude, the loaded and unloaded *Q* factors can be extracted with substantially higher accuracy. Such a setup would also permit measurement of the coupling coefficient *k* from the splitting of resonance peaks in a coupled resonator system. Replacing alligator clips with soldered or crimped connections and trimming the extra 75 cm of wire would eliminate contact resistance variability and unmodelled AC resistance contributions.

6.9.2 Ferrite-Related Uncertainties

The ferrite tiles used in the fully assembled couplers have a specified initial permeability $\mu_i = 3000 \pm 25\%$. The Maxwell simulations assumed the nominal value, but the measured self- and mutual inductances for couplings *A-II* and *B* were approximately 10–15% lower than predicted – consistent with a ferrite batch at the lower end of the tolerance band. This explains the mutual inductance errors of up to 10% observed for these couplings. For coupling *A-I*, which uses a

different ferrite layout with less coverage, the sensitivity to permeability variations is reduced, and the errors are much smaller (below 1%).

The key implication is that while the simulations capture the behaviour correctly, the quantitative agreement depends on an unknown parameter. Without characterising the actual permeability of the ferrite batch, the absolute accuracy of the ferrite-loaded simulations remains uncertain. However, the design ranking is preserved regardless of the absolute values, confirming that the optimisation methodology is robust to batch-to-batch ferrite variability.

Future direction: A simple inductance measurement on a standard test fixture would provide a batch-specific value of the ferrite's initial permeability. Repeating the Stage 2–4 simulations with this measured value would improve quantitative agreement – a low-cost, high-impact improvement for future work.

6.9.3 Limitations of the Test Jig

Mechanical slack in the 3D-printed jig led to positioning errors, particularly at extreme misalignment offsets. This is the most plausible explanation for the counter-intuitive finding that mutual inductance errors are smaller at extreme misalignments than at the aligned position. The mechanical slack effectively reduces the misalignment at extreme offsets, partially compensating for the lower ferrite permeability. This limitation affects the accuracy of positioning and, consequently, the measured mutual inductance values.

In addition to mechanical slack, the thermal properties of the PLA components – with a glass transition temperature of only 55–65°C – represent a significant barrier to full-power validation. The thermal simulations for coupling *A-I* predicted a coil temperature of 52.5°C, already close to the lower end of this range. While PLA is adequate for the signal-level measurements performed here, high-power testing would risk softening and deformation of the printed components.

Future direction: A redesigned test jig incorporating precision linear stages with digital encoders would eliminate mechanical slack and improve repeatability. Additionally, the 3D-printed components should be reprinted using a more adequate filament (PETG) to provide adequate thermal margin for full-power tests.

6.9.4 Scope Limitations

Beyond the specific measurement challenges, the experimental validation was limited in scope in two important ways. First, the validation was conducted at signal level, not at the rated 3.7 kW. While this is sufficient for validating inductance and coupling measurements, it does not validate the thermal predictions of Stage 5 or the ferrite saturation analysis of Stage 4. Second, couplings *C* and *D* – which use the finer 1200×AWG38 Litz wire – remain as validated simulation designs awaiting physical implementation. These designs represent a higher-performance trajectory that could not be realised within the project timeline due to the unavailability of the finer Litz wire.

Future direction: Full-power testing at 3.7 kW is essential to validate thermal predictions and saturation analysis, requiring a test jig with adequate thermal management and components rated

for full power. Couplings *C* and *D* should be wound using the same modular corner-holder design, with components reprinted in high-temperature filament, and full-power tests performed to validate the higher-performance trajectory predicted by simulations.

6.10 Chapter Summary

This chapter has presented the practical implementation and experimental validation of couplings *A-I*, *A-II*, and *B*, providing the ultimate test of the design methodology. The mechanical test jig, constructed from non-magnetic materials with 3D-printed PLA components, successfully enabled three-axis positioning over the SAE J2954 Z1 envelope. The modular, reconfigurable corner-holder design proved effective in reproducing the intended coil geometries, as confirmed by air-core inductance measurements (errors < 3.2%).

The experimental results validated the fundamental design ranking: coupling *A-II* consistently outperformed *A-I* across all measured positions, with efficiency improvements of 0.3–0.5 percentage points in both simulation and measurement. This ranking consistency is the fundamental pass/fail criterion for the pre-design pipeline, demonstrating that the toolchain correctly identifies relative performance even when absolute errors remain.

Despite the limitations discussed in Section 6.9 the experimental results are sufficient to validate the fundamental capability of the design methodology. The self-inductance of air-core coils is predicted to within 3.2%; the self-inductance of ferrite-loaded coils is predicted to within 11.5% (with the discrepancy attributable to ferrite tolerance); the mutual inductance trends with misalignment are correctly captured; and the maximum efficiency remains within 0.5 percentage points of the simulated values for the air-core case.

The toolchain correctly identifies relative performance even when absolute errors remain, demonstrating that the methodology is fit for its intended purpose: screening and ranking candidate designs prior to FEA refinement and hardware realisation. The identified limitations provide clear directions for improvement, and a higher-level synthesis of these limitations and future directions, contextualised within the analytical toolchain and complete design methodology, is provided in Chapter 7.

Chapter 7

Conclusions and Future Work

This final chapter synthesises the contributions of this thesis, reflects on the broader implications of the work, and outlines a roadmap for future research. The thesis has developed a complete design pipeline for inductive power transfer (IPT) magnetic couplers, demonstrating how a fast analytical pre-filter can be systematically integrated with high-fidelity finite-element simulations and experimental validation to enable efficient and reliable coupler design. The central contribution is a five-stage methodology that bridges the gap between rapid analytical estimation and slow, high-fidelity FEA simulation, focusing computational resources on the most promising candidates.

The three sections that follow address each pillar of the work – the analytical toolchain, the complete design methodology, and the experimental implementation – as a broad synthesis of what was achieved and where the work points next. Detailed discussions of specific limitations and future directions are provided in the respective chapters: Section 4.7 for the analytical toolchain, Section 5.9 for the complete design methodology, and Section 6.9 for the experimental implementation.

7.1 The Analytical Toolchain

The MATLAB toolchain developed in Chapter 4 establishes a new capability for rapid, physics-informed pre-design of IPT magnetic couplers. By combining the discretised Neumann integral with a per-segment Biot–Savart proximity loss model, the toolchain achieves a favourable balance between computational speed and physical accuracy. Its ability to evaluate coupling statistics over the full SAE J2954 Z1 envelope – or any arbitrary 3D region – moves the field beyond the single-position optimisation that has dominated much of the literature. The validation against Ansys Maxwell FEA confirms that the toolchain preserves design ranking, enabling it to function as a reliable pre-filter that focuses expensive computational resources on genuinely promising candidates.

The toolchain’s diagnostic capability was demonstrated through the proximity-loss comparison between coarse and fine Litz wire constructions, revealing how strand diameter directly influences AC resistance and quality factor. This insight informed the selection of wire constructions for the five candidate designs and illustrates the broader value of physics-informed loss modelling in

guiding design decisions. The toolchain is not intended to replace FEA, but rather to complement it by enabling rapid exploration of the design space and identifying candidates worthy of detailed simulation.

The limitations of the toolchain – its air-core scope, the absence of ferrite and shielding models, the heuristic nature of the optimisation, and the finite set of supported geometries – are acknowledged and discussed in detail in Section 4.7. These limitations do not undermine the toolchain’s utility for its intended purpose, but they define its scope and point toward clear directions for future development, including analytical ferrite models, multi-receiver optimisation, and extended geometry support.

7.2 The Complete Design Methodology

The five-stage pipeline presented in Chapter 5 demonstrates how the analytical toolchain can be integrated with FEA validation, ferrite and shielding integration, electrical worst-case assessment, and thermal simulation to produce complete, physically validated coupler designs. The methodology was applied to five candidate couplings, establishing a coherent set of designs whose electromagnetic, electrical, and thermal behaviour is fully characterised across the SAE J2954 Z1 misalignment envelope.

The comparison between coupling A-I and coupling A-II served as a proof-of-concept for the necessity of accurate loss modelling, demonstrating that a toolchain blind to proximity losses converges on a suboptimal “touching turns” design. The spatially robust ferrite layout selection, based on performance averaged over the full misalignment envelope, ensured that the chosen configurations perform reliably under realistic parking conditions rather than only at the centred position. The design-closure framework – with built-in validation checkpoints that retrospectively confirm initial assumptions – provides a systematic basis for iteration and ensures that design decisions are verified before proceeding to hardware fabrication.

The cost-performance metric $\Delta\eta/\mathcal{C}$ introduced in Section 5.5.7 provides a rational basis for ferrite layout selection that explicitly balances performance gain against material cost, addressing the economic constraints of real-world engineering. The resonance detuning sensitivity analysis further emphasised that a fixed-frequency operating strategy is not robust against typical component tolerances, underscoring the need for adaptive control strategies.

The limitations of the pipeline – including the standardisation of ferrite layouts, the manual nature of the FEA setup, the assumptions in the thermal simulations, the dependency on the SS compensation topology, and the absence of tolerance analysis – are discussed in detail in Section 5.9. These limitations stem from pragmatic choices made to keep the five-coupling comparison tractable, and they point toward clear directions for future development, including per-coupling ferrite optimisation, automated FEA integration, and system-level design extension.

7.3 Experimental Implementation

The practical implementation and experimental validation of couplings *A-I*, *A-II*, and *B*, presented in Chapter 6, provided the ultimate test of the design methodology. The mechanical test jig, constructed from non-magnetic materials with 3D-printed PLA components, successfully enabled three-axis positioning over the SAE J2954 Z1 envelope, and the modular, reconfigurable corner-holder design proved effective in reproducing the intended coil geometries.

The experimental results confirmed the fundamental design ranking established in simulation: coupling *A-II* consistently outperformed *A-I* across all measured positions. This ranking consistency is the fundamental pass/fail criterion for the pre-design pipeline, demonstrating that the toolchain correctly identifies relative performance even when absolute errors remain. The air-core inductance measurements agreed with simulations to within 3.2 %, confirming that the winding process and the 3D-printed holders successfully reproduced the designed geometries.

The experimental validation also revealed the practical challenges of characterising high- Q IPT couplers. The phase sensitivity of one-port impedance measurements, the variability of ferrite permeability, and the mechanical limitations of 3D-printed test fixtures are constraints that future work must address. The candid documentation of these challenges is itself a contribution, providing a realistic account of the gap between simulation and measurement that is rarely discussed in the literature.

The limitations of the experimental work – including measurement-induced errors, ferrite-related uncertainties, mechanical and thermal limitations of the test jig, and scope limitations – are discussed in detail in Section 6.9. These limitations point toward clear directions for future development, including forced-resonance measurement setups, ferrite permeability characterisation, redesigned test jigs with precision stages, and full-power validation of the remaining designs.

7.4 Publication and Dissemination

A key outcome of this work is the dissemination of the analytical toolchain to the broader research community. An abstract entitled “A MATLAB-Based Design and Optimisation Tool for Wireless Power Transfer Magnetic Couplers” has been accepted for poster presentation at the 10th E-Mobility Power System Integration Symposium, taking place on 28–29 September 2026 in Porto, Portugal.

The paper focuses on the MATLAB toolchain developed in this thesis (Chapter 4) and presents it as a standalone resource for the rapid, air-core pre-design of IPT magnetic couplers. The toolchain’s ability to compute self-inductance, mutual inductance, and AC resistance using the discretised Neumann integral and a per-segment proximity loss model, combined with its capacity to evaluate coupling statistics over arbitrary 3D misalignment envelopes, makes it applicable to a wide range of IPT use cases beyond the SAE J2954 framework demonstrated in this work.

The primary motivation for this publication is to make the toolchain available to the scientific community, inviting feedback on its utility, limitations, and potential extensions. By presenting

the tool in a forum dedicated to e-mobility power systems, the work aims to reach researchers and practitioners who may benefit from a fast, open, and flexible pre-design filter for IPT couplers. The feedback gathered from the community will inform future refinements of the toolchain and guide its evolution towards broader applicability.

The paper is currently in preparation and will be submitted prior to the symposium. The acceptance of the abstract confirms the relevance of the toolchain to ongoing research in wireless power transfer and e-mobility, and represents an important step in disseminating the contributions of this thesis beyond the academic environment where it was developed.

7.5 Final Closing Statement

This thesis has presented a complete, standards-compliant design pipeline for IPT magnetic couplers that effectively bridges the gap between fast analytical estimation and slow, high-fidelity FEA validation. The five-stage methodology – comprising analytical pre-filtering, air-core FEA validation, ferrite and shielding integration, electrical worst-case assessment, and thermal simulation – provides a systematic framework that enables engineers to navigate the complex, multi-dimensional design space with confidence. The toolchain’s ability to diagnose proximity losses and to maintain design ranking across all stages has been validated through comprehensive simulation and experimental testing.

The practical implementation of couplings *A-I*, *A-II*, and *B* demonstrated that the methodology translates directly into physical hardware with measured performance aligned with predictions. The remaining designs (*C* and *D*) offer a pathway to even higher performance through the use of finer Litz wire. The identified limitations – from the toolchain’s air-core scope through the pipeline’s standardised ferrite layouts and fixed SS topology to the experimental jig’s measurement phase sensitivity and thermal constraints – provide clear directions for future refinement.

The experimental validation, while confirming the fundamental design ranking, also revealed the practical challenges of characterising high- Q IPT couplers. The phase sensitivity of one-port impedance measurements, the variability of ferrite permeability, and the mechanical limitations of 3D-printed test fixtures are all constraints that future work must address. The candid documentation of these challenges is itself a contribution, providing a realistic account of the gap between simulation and measurement that is rarely discussed in the literature.

The dissemination of the analytical toolchain through the 10th E-Mobility Power System Integration Symposium (Section 7.4) marks an important step in making this work accessible to the broader research community. The feedback and engagement from this effort will inform future refinements and extensions of the methodology.

As wireless power transfer continues its transition from academic research to widespread industrial deployment, the demand for efficient, reliable, and cost-effective design tools will only grow. The methodology presented in this thesis offers a robust foundation for meeting that demand, and the future research directions identified throughout this work provide a roadmap for

its continued evolution towards a fully integrated, multi-objective, and experimentally validated design framework.

References

- [1] Osama Abuajwa et al. “Comprehensive review of wireless power transfer systems for electric vehicle charging applications”. In: *Discover Applied Sciences* 7.10 (2025). DOI: [10.1007/s42452-025-07738-z](https://doi.org/10.1007/s42452-025-07738-z).
- [2] All About Circuits. *4-wire Resistance Measurement*. Last accessed: 2026-06. URL: <https://www.allaboutcircuits.com/textbook/direct-current/chpt-8/kelvin-resistance-measurement/>.
- [3] A. Alphones and Prasad Jayathurathnage. “Review on Wireless Power Transfer Technology (invited paper)”. In: *IEEE Asia Pacific Microwave Conference (APMC)*. Sept. 2017.
- [4] Ansys, Inc. *Icepak 2025 R2 Tutorials*. Last accessed: 2026-06-15. URL: <https://ansyshelp.ansys.com/public/account/secured?returnurl=/Views/Secured/Electronics/v252/en/Subsystems/Icepak/Icepak.htm>.
- [5] Ansys, Inc. *Maxwell 2025 R2 Tutorials*. Last accessed: 2026-06-15. URL: https://ansyshelp.ansys.com/public/account/secured?returnurl=/Views/Secured/prod_page.html?pn=Maxwell&pid=Maxwell&prodver=25.2&lang=en.
- [6] Ilyas Aslan and Atiye Hulya Obdan. “Design Optimization of Circular Coils in Wireless Electric Vehicle Charging Systems”. In: *Journal of Magnetism* 29.4 (2024), pp. 358–368. DOI: [10.4283/JMAG.2024.29.4.358](https://doi.org/10.4283/JMAG.2024.29.4.358).
- [7] Charles Marfo Awuah et al. “Novel coil design and analysis for high-power wireless power transfer with enhanced Q-factor”. In: *Scientific Reports* 13.1 (2023). DOI: [10.1038/s41598-023-31389-y](https://doi.org/10.1038/s41598-023-31389-y).
- [8] Slobodan Babic et al. “Mutual inductance calculation between circular filaments arbitrarily positioned in space: Alternative to grover’s formula”. In: *IEEE Transactions on Magnetics* 46.9 (2010), pp. 3591–3600. DOI: [10.1109/TMAG.2010.2047651](https://doi.org/10.1109/TMAG.2010.2047651).
- [9] Nguyen Xuan Bac, D. Mahinda Vilathgamuwa, and Udaya K. Madawala. “A SiC-based matrix converter topology for inductive power transfer system”. In: *IEEE Transactions on Power Electronics* 29.8 (2014), pp. 4029–4038. DOI: [10.1109/TPEL.2013.2291434](https://doi.org/10.1109/TPEL.2013.2291434).
- [10] Mohamed Bensetti et al. “Parametric Optimization of Ferrite Structure Used for Dynamic Wireless Power Transfer for 3 kW Electric Vehicle”. In: *Energies* 16.14 (2023). DOI: [10.3390/en16145439](https://doi.org/10.3390/en16145439).
- [11] Mickel Budhia, Grant A. Covic, and John T. Boys. “Design and optimization of circular magnetic structures for lumped inductive power transfer systems”. In: *IEEE Transactions on Power Electronics* 26.11 (2011), pp. 3096–3108. DOI: [10.1109/TPEL.2011.2143730](https://doi.org/10.1109/TPEL.2011.2143730).
- [12] Soham Chatterjee et al. “Design Optimisation for an Efficient Wireless Power Transfer System for Electric Vehicles”. In: *Energy Procedia*. Vol. 117. 2017, pp. 1015–1023. DOI: [10.1016/j.egypro.2017.05.223](https://doi.org/10.1016/j.egypro.2017.05.223).

- [13] Tianyuan Chen et al. *An overview of power loss calculation methods for high-frequency litz wires*. July 2024. DOI: [10.3233/JAE-230171](https://doi.org/10.3233/JAE-230171).
- [14] Components 101. *How Does Quality Factor Impact the Performance of an Inductor?* Last accessed: 2026-06-05. URL: <https://components101.com/articles/how-does-quality-factor-impact-the-performance-of-an-inductor>.
- [15] Fabio Corti et al. “MatWPT: A MATLAB Tool for Wireless Power Transfer Systems Design Optimization”. In: *IEEE Open Journal of Industry Applications* 6 (2025), pp. 415–426. DOI: [10.1109/OJIA.2025.3580262](https://doi.org/10.1109/OJIA.2025.3580262).
- [16] Grant A. Covic and John T. Boys. “Inductive power transfer”. In: *Proceedings of the IEEE* 101.6 (2013), pp. 1276–1289. DOI: [10.1109/JPROC.2013.2244536](https://doi.org/10.1109/JPROC.2013.2244536).
- [17] Kalina Detka and Krzysztof Górecki. “Wireless Power Transfer—A Review”. In: *Energies* 15.19 (2022). DOI: [10.3390/en15197236](https://doi.org/10.3390/en15197236).
- [18] Andrew Foote and Omer C Onar. *A Review of High-Power Wireless Power Transfer*. Tech. rep. 2017.
- [19] Tarek S. Hassan, Mohamed M. Zakaria Moustafa, and Nabil H. Abbasy. “High-Efficiency Wireless Power Transfer for Electric Vehicles: A Coil Geometry and PSO-Based Optimization Framework”. In: *IEEE Access* 13 (2025), pp. 132952–132972. DOI: [10.1109/ACCESS.2025.3592070](https://doi.org/10.1109/ACCESS.2025.3592070).
- [20] Dejana Herceg et al. “Double-Layer Coils Design for 11 kW Wireless Power Transfer”. In: *Electronics (Switzerland)* 13.3 (Feb. 2024). DOI: [10.3390/electronics13030547](https://doi.org/10.3390/electronics13030547).
- [21] Iftikhar Hussain and Dong Kyun Woo. “Inductance Calculation of Single-Layer Planar Spiral Coil”. In: *Electronics (Switzerland)* 11.5 (2022). DOI: [10.3390/electronics11050750](https://doi.org/10.3390/electronics11050750).
- [22] Rahulkumar J. et al. “A review on resonant inductive coupling pad design for wireless electric vehicle charging application”. In: *Energy Reports* 10 (2023). DOI: [10.1016/j.egy.2023.08.067](https://doi.org/10.1016/j.egy.2023.08.067).
- [23] Soukaina Jaafari et al. “Optimization of circular coils with ferrite boxes for enhanced efficiency in wireless power transfer for electric vehicles”. In: *Green Energy and Intelligent Transportation* 4.2 (2025). DOI: [10.1016/j.geits.2024.100195](https://doi.org/10.1016/j.geits.2024.100195).
- [24] Kafeel Ahmed Kalwar et al. “Coil design for high misalignment tolerant inductive power transfer system for EV charging”. In: *Energies* 9.11 (2016). DOI: [10.3390/en9110937](https://doi.org/10.3390/en9110937).
- [25] Minh Kim et al. *Low Frequency Electromagnetic Compatibility of Wirelessly Powered Electric Vehicles*. Tech. rep. 2014.
- [26] Yusuf Kosesoy et al. *Design and Implementation of a 1-kW Wireless Power Transfer System for EV Charging*. Tech. rep. 2019.
- [27] Damien Lemaitre et al. *Efficiency-Oriented Design of Litz Wire for Several kW Power Experimented on 20 kW Prototype. PCIM conference*. Tech. rep. 2021. URL: <https://hal.science/hal-04733625v1>.
- [28] Siqi Li and Chunting Chris Mi. “Wireless power transfer for electric vehicle applications”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 3.1 (2015), pp. 4–17. DOI: [10.1109/JESTPE.2014.2319453](https://doi.org/10.1109/JESTPE.2014.2319453).
- [29] Jiajing Liang. “Benchmarking of cutting-edge wireless power transfer for electric vehicles”. In: *Applied and Computational Engineering* 10.1 (2023), pp. 240–249. DOI: [10.54254/2755-2721/10/20230184](https://doi.org/10.54254/2755-2721/10/20230184).

- [30] Rüdiger Matthes, Jürgen H. Bernhardt, and Alastair F. McKinlay. *Guidelines on limiting exposure to non-ionizing radiation*. International Commission on Non-Ionizing Radiation Protection, 1998.
- [31] Elisa de Melo Henriques and Sascha Stegen. “Concave Ferrite Core for Wireless Power Transfer (WPT)”. In: *Energies* 16.12 (2023). DOI: [10.3390/en16124553](https://doi.org/10.3390/en16124553).
- [32] Ahmed A.S. Mohamed, Siguang An, and Osama Mohammed. “Coil Design Optimization of Power Pad in IPT System for Electric Vehicle Applications”. In: *IEEE Transactions on Magnetics* 54.4 (2018). DOI: [10.1109/TMAG.2017.2784381](https://doi.org/10.1109/TMAG.2017.2784381).
- [33] Ahmed A.S. Mohamed et al. “A comprehensive overview of inductive pad in electric vehicles stationary charging”. In: *Applied Energy* 262 (2020). DOI: [10.1016/j.apenergy.2020.114584](https://doi.org/10.1016/j.apenergy.2020.114584).
- [34] Chandrasekharan Nataraj et al. “Analysis of mutual inductance and coupling factor of inductively coupled coils for wireless electricity”. In: *ARPJ Journal of Engineering and Applied Sciences* 12.13 (2017). URL: www.arpnjournals.com.
- [35] Minh T. Nguyen et al. “Electromagnetic field based WPT technologies for UAVs: A comprehensive survey”. In: *Electronics (Switzerland)* 9.3 (2020). DOI: [10.3390/electronics9030461](https://doi.org/10.3390/electronics9030461).
- [36] International Commission on Non-Ionizing Radiation Protection. “Guidelines for limiting exposure to time-varying electric and magnetic fields (1 Hz TO 100 kHz)”. In: *Health Physics* 99.6 (2010), pp. 818–836. DOI: [10.1097/HP.0b013e3181f06c86](https://doi.org/10.1097/HP.0b013e3181f06c86).
- [37] Daniel Ongayo and Moin Hanif. “Comparison of circular and rectangular coil transformer parameters for wireless Power Transfer based on Finite Element Analysis”. In: *2015 IEEE 13th Brazilian Power Electronics Conference and 1st Southern Power Electronics Conference, COBEP/SPEC 2016*. 2015. DOI: [10.1109/COBEP.2015.7420222](https://doi.org/10.1109/COBEP.2015.7420222).
- [38] Pedro Gelati Pascoal. *Projeto do Acoplamento Indutivo de um Sistema de Transferência de Energia sem Fio para a Recarga de Veículos Elétricos*. Tech. rep. Universidade Federal de Santa Maria, Santa Maria, RS, Brazil, 2023.
- [39] Pedro Gelati Pascoal et al. “Projeto e Simulação das Bobinas de um Sistema de Transferência de Energia Sem Fio para Aplicação na Recarga de Veículos Elétricos”. In: *Proceedings do XXIV Congresso Brasileiro de Automática*. 2022. DOI: [10.20906/cba2022/3547](https://doi.org/10.20906/cba2022/3547).
- [40] Ali Pashaei et al. “A novel hybrid coil design and implementation for wireless power transfer systems”. In: *Science and Technology for Energy Transition (STET)* 79 (2024). DOI: [10.2516/stet/2024082](https://doi.org/10.2516/stet/2024082).
- [41] Nataša Prosen, Jure Domajnko, and Miro Milanovič. “Wireless power transfer using double DD coils”. In: *Electronics (Switzerland)* 10.20 (Oct. 2021). DOI: [10.3390/electronics10202528](https://doi.org/10.3390/electronics10202528).
- [42] Jaewon Rhee et al. “Selection of Ferrite Depending on Permeability and Weight to Enhance Power Transfer Efficiency in Low-Power Wireless Power Transfer Systems”. In: *Energies* 17.15 (2024). DOI: [10.3390/en17153816](https://doi.org/10.3390/en17153816).
- [43] *SAE J2954: Wireless Power Transfer for Light-Duty Plug-In/Electric Vehicles and Alignment Methodology*. Standard J2954. SAE International, 2024.
- [44] Viktor Shevchenko et al. “Coil Design for Wireless Power Transfer with Series-Parallel Compensation”. In: *IEEE 2nd Ukraine Conference on Electrical and Computer Engineering Lviv*. 2019. DOI: [10.1109/UKRCON.2019.8879877](https://doi.org/10.1109/UKRCON.2019.8879877).

- [45] Zhiyuan Shi et al. *High Efficiency Wireless Power Transfer System for Medium Power*. Tech. rep. 2016.
- [46] Nahed H. Solouma et al. “An Efficient Design of Inductive Transmitter and Receiver Coils for Wireless Power Transmission”. In: *Electronics (Switzerland)* 12.3 (Feb. 2023). DOI: [10.3390/electronics12030564](https://doi.org/10.3390/electronics12030564).
- [47] Linlin Tan et al. “Comparative analysis and design of the shielding techniques in WPT systems for charging EVs”. In: *Energies* 12.11 (2019). DOI: [10.3390/en12112115](https://doi.org/10.3390/en12112115).
- [48] United Nations. *The Sustainable Development Goals*. <https://sdgs.un.org/goals>. Accessed: 2026-06-20. 2025.
- [49] Ravikiran Vaka and Ritesh Kumar Keshri. “Review on contactless power transfer for electric vehicle charging”. In: *Energies* 10.5 (2017). DOI: [10.3390/en10050636](https://doi.org/10.3390/en10050636).
- [50] Leonardo A. Brum Viera. *Análise e Desenvolvimento de um Sistema de Transferência de Energia sem Fio Aplicado na Recarga de Veículos Elétricos*. Tech. rep. Universidade Federal de Santa Maria, Santa Maria, RS, Brazil, 2022.
- [51] Leonardo A. Brum Viera, Pedro Gelati Pascoal, and Cassiano Rech. “Optimized Design Methodology and Maximum Efficiency Tracking Algorithm for Static IPT Chargers in Electric Vehicles”. In: *Eletrônica de Potência* 29 (2024), e202424. DOI: [10.18618/rep.2005.2.023027](https://doi.org/10.18618/rep.2005.2.023027).
- [52] Leonardo A. Brum Viera, Pedro Gelati Pascoal, and Cassiano Rech. “Optimized Design Methodology in Inductive Power Transfer Systems Applied to Electric Vehicle Charging”. In: *COBEP 2023 - 17th Brazilian Power Electronics Conference and SPEC 2023 - 8th IEEE Southern Power Electronics Conference, Proceedings*. 2023. DOI: [10.1109/SPEC56436.2023.10407670](https://doi.org/10.1109/SPEC56436.2023.10407670).
- [53] Leonardo A. Brum Viera et al. “Análise e Projeto de um Sistema de Recarga de Veículos Elétricos através da Transferência de Energia sem Fio”. In: *Seminar on Power Electronics and Control - SEPOC*. 2021. DOI: [10.53316/sepoc2021.035](https://doi.org/10.53316/sepoc2021.035).
- [54] Leonardo Antonio Brum Viera. *CONTRIBUIÇÕES AO DESENVOLVIMENTO DE UM TRANSMISSOR INTEROPERÁVEL PARA SISTEMAS DE RECARGA SEM FIO ESTÁTICOS DE VEÍCULOS ELÉTRICOS*. Tech. rep. Universidade Federal de Santa Maria, Santa Maria, RS, Brazil, 2025.
- [55] De'an Wang et al. “Modern Advances in Magnetic Materials of Wireless Power Transfer Systems: A Review and New Perspectives”. In: *Nanomaterials* 12.20 (2022). DOI: [10.3390/nano12203662](https://doi.org/10.3390/nano12203662).
- [56] Guo Wei et al. “An Automatic Coil Design Method with Modified AC Resistance Evaluation for Achieving Maximum Coil-Coil Efficiency in WPT Systems”. In: *IEEE Transactions on Power Electronics* 35.6 (2020), pp. 6114–6126. DOI: [10.1109/TPEL.2019.2952120](https://doi.org/10.1109/TPEL.2019.2952120).
- [57] Rafal P. Wojda and Marian K. Kazimierczuk. “Winding Resistance and Power Loss of Inductors with Litz and Solid-Round Wires”. In: *IEEE Transactions on Industry Applications* 54.4 (July 2018), pp. 3548–3557. DOI: [10.1109/TIA.2018.2821647](https://doi.org/10.1109/TIA.2018.2821647).
- [58] Yang Yang, Jinlong Cui, and Xin Cui. “Design and analysis of magnetic coils for optimizing the coupling coefficient in an electric vehicle wireless power transfer system”. In: *Energies* 13.6 (2020). DOI: [10.3390/en13164143](https://doi.org/10.3390/en13164143).
- [59] Tuba Yilmaz et al. “Multi-Objective Optimization of Circular Magnetic Couplers for Wireless Power Transfer Applications”. In: *IEEE Transactions on Magnetics* 53.8 (Aug. 2017). DOI: [10.1109/TMAG.2017.2692218](https://doi.org/10.1109/TMAG.2017.2692218).

- [60] Jian Yin et al. “Front-End Monitoring of the Mutual Inductance and Load Resistance in a Series-Series Compensated Wireless Power Transfer System”. In: *IEEE Transactions on Power Electronics* 31.10 (Oct. 2016), pp. 7339–7352. DOI: [10.1109/TPEL.2015.2509962](https://doi.org/10.1109/TPEL.2015.2509962).
- [61] Mohammad Zamani et al. “A review of inductive power transfer for electric vehicles”. In: *IEEE* (2019).
- [62] Martin Zavřel and Vladimír Kindl. “Comparison of Various Configuration of Wireless Power Transfer System”. In: *Transactions on Electrical Engineering* 7.2 (2020), pp. 32–36. DOI: [10.14311/tee.2018.2.032](https://doi.org/10.14311/tee.2018.2.032).
- [63] Zhen Zhang et al. “Wireless Power Transfer - An Overview”. In: *IEEE Transactions on Industrial Electronics* 66.2 (2019), pp. 1044–1058. DOI: [10.1109/TIE.2018.2835378](https://doi.org/10.1109/TIE.2018.2835378).

Appendix A

Validation Couplings Configurations

This appendix summarises the complete set of Tx–Rx coil configurations used for the coupling analysis in Chapter 4. All configurations assume a fixed evaluation region defined by the `REGION` variable (see Section 4.6). Conductor types `LITZ_A`, `LITZ_B`, `LITZ_C`, and `LITZ_D` are defined in Table 4.13.

Note: In the following tables, all dimensions, corner radii, turn spacings, and layer spacings are in millimetres (mm).

A.1 T1: Single-layer Spiral – Single-layer Spiral

Table A.1: T1 Configurations – geometric variation.

	Topology	Size	Corner	Turns	Turn Spacing	Layers	Layer Spacing	Conductor
Label	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx
T1-1	Spiral	300x300	0	8	2	1	–	LITZ_A
T1-2	Spiral	300x300	60	8	2	1	–	LITZ_A
T1-3	Spiral	300x300	150	5	2	1	–	LITZ_A
T1-4	Spiral	400x250	0	8	2	1	–	LITZ_A
T1-5	Spiral	400x250	60	8	2	1	–	LITZ_A
T1-6	Spiral	300x300	60	12	0	1	–	LITZ_A
T1-7	Spiral	300x300	60	5	5	1	–	LITZ_A
T1-8	Spiral	500x500	100	10	3	1	–	LITZ_A

A.2 T2: Single-layer Spiral – Single-layer Spiral

Table A.2: T2 Configurations – size variation.

Label	Topology	Size		Corner		Turns		Turn Spacing	Layers	Layer Spacing	Conductor
	Tx/Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx
T2-1	Spiral	300x300	200x200	60	40	8	6	2	1	–	LITZ_A
T2-2	Spiral	400x400	300x300	80	60	10	8	2	1	–	LITZ_A
T2-3	Spiral	500x500	200x200	100	40	10	6	3	1	–	LITZ_A
T2-4	Spiral	400x250	200x200	60	40	8	6	2	1	–	LITZ_A

A.3 T3: Single-layer Spiral – Single-layer Spiral

Table A.3: T3 Configurations – conductor variation.

Label	Topology	Size	Corner	Turns	Turn Spacing	Layers	Layer Spacing	Conductor	
	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx	Rx
T3-1	Spiral	300x300	60	8	2	1	–	LITZ_B	LITZ_A
T3-2	Spiral	300x300	60	8	2	1	–	LITZ_C	LITZ_A
T3-3	Spiral	300x300	60	8	2	1	–	LITZ_D	LITZ_A
T3-4	Spiral	300x300	60	8	2	1	–	LITZ_B	LITZ_B
T3-5	Spiral	300x300	60	8	2	1	–	LITZ_C	LITZ_D

A.4 T4: Single-layer DD – Single-layer Spiral

Table A.4: T4 Configurations – mixed topology.

Label	Topology		Size		Corner		Turns		Turn Spacing	Layers	Layer Spacing	Conductor
	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx
T4-1	DD	Spiral	200x300	300x300	60	60	8	8	2	1	–	LITZ_A
T4-2	DD	Spiral	200x300	200x200	60	40	8	6	2	1	–	LITZ_A
T4-3	DD	Spiral	250x400	300x300	100	60	10	8	2	1	–	LITZ_A

A.5 T5: Single-layer DD – Single-layer DD

Table A.5: T5 Configurations – both coils Double-D.

Label	Topology	Size		Corner		Turns		Turn Spacing	Layers		Layer Spacing	Conductor
	Tx/Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx
T5-1	DD	200x300	200x300	60	60	8	8	2	1	–		LITZ_A
T5-2	DD	200x300	150x200	60	40	8	6	2	1	–		LITZ_A

A.6 T6: Multi-layer Spiral – Single-layer Spiral

Table A.6: T6 Configurations – multi-layer effects.

Label	Topology	Size	Corner	Turns	Turn Spacing	Layers		Layer Spacing		Conductor
	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx	Rx	Tx	Rx	Tx/Rx
T6-1	Spiral	300x300	60	8	2	2	1	0	–	LITZ_A
T6-2	Spiral	300x300	60	8	2	2	1	8	–	LITZ_A
T6-3	Spiral	300x300	60	8	2	2	1	15	–	LITZ_A

A.7 T7: Multi-layer Spiral – Multi-layer Spiral

Table A.7: T7 Configurations – both coils multi-layer Spiral.

Label	Topology	Size	Corner	Turns	Turn Spacing	Layers		Layer Spacing		Conductor
	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx
T7-1	Spiral	300x300	60	8	2	2		8		LITZ_A
T7-2	Spiral	300x300	60	8	2	2		15		LITZ_A

A.8 T8: Multi-layer DD – Single-layer Spiral

Table A.8: T8 Configurations – mixed topology.

Label	Topology		Size		Corner	Turns	Turn Spacing	Layers		Layer Spacing		Conductor
	Tx	Rx	Tx	Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx	Rx	Tx	Rx	Tx/Rx
T8-1	DD	Spiral	200x300	300x300	60	8	2	2	1	0	–	LITZ_A
T8-2	DD	Spiral	200x300	300x300	60	8	2	2	1	8	–	LITZ_A
T8-3	DD	Spiral	200x300	300x300	60	8	2	2	1	15	–	LITZ_A

A.9 T9: Multi-layer DD – Multi-layer DD

Table A.9: T9 Configurations – both coils multi-layer Double-D.

	Topology	Size	Corner	Turns	Turn Spacing	Layers	Layer Spacing	Conductor
Label	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx/Rx
T9-1	DD	200x300	60	8	2	2	8	LITZ_A
T9-2	DD	200x300	60	8	2	2	15	LITZ_A

A.10 T10: Mixed conductor and topology pairs

Table A.10: T10 Configurations – mixed geometries and conductors.

	Topology		Size		Corner	Turns	Turn Spacing	Layers		Layer Spacing		Conductor	
Label	Tx	Rx	Tx	Rx	Tx/Rx	Tx/Rx	Tx/Rx	Tx	Rx	Tx	Rx	Tx	Rx
T10-1	Spiral	Spiral	300x300	300x300	60	8	2	2	1	8	–	LITZ_A	LITZ_B
T10-2	Spiral	Spiral	300x300	300x300	60	8	2	2	1	8	–	LITZ_B	LITZ_C
T10-3	DD	Spiral	200x300	300x300	60	8	2	2	1	8	–	LITZ_A	LITZ_D
T10-4	DD	Spiral	200x300	300x300	60	8	2	2	2	8	8	LITZ_C	LITZ_D

Appendix B

Ferrite Layouts and Associated Metrics

This appendix presents the absolute electrical parameters and coupling properties for all ferrite layout configurations analysed in Chapter 5. For each coupling and ferrite arrangement, the tables report the coil resistance R , self-inductance L , quality factor Q , spatially averaged mutual inductance M_{avg} , coupling coefficient k_{avg} , maximum efficiency η_{avg} , ferrite tile counts (#100 and #150), total ferrite cost, and the cost-performance metric $\Delta\eta/\mathcal{C}$.

Note: All ferrite counts and costs refer to the full coupling (both Tx and Rx coils combined), and all values are spatial averages across the 60-position evaluation grid. The cost-performance metric is defined in Section 5.5.7.

B.1 Coupling A - I

Table B.1: Coupling A - I : absolute electrical values and cost for each ferrite layout.

Config.	R	L	Q	M	k	η_{max}	Ferrite count		Cost	$\Delta\eta/\text{Cost}$
	(m Ω)	(μH)	(–)	(μH)	(–)	(%)	#100	#150	(€)	($\times 10^6$ % per €)
Air-core	97.95	130.89	713.7	30.44	0.2326	98.77	–	–	0.00	–
Type X	138.23	158.76	613.4	31.70	0.1996	98.29	0	24	581.04	–826.0

B.2 Coupling A-II

Table B.2: Coupling A-II: absolute electrical values and cost for each ferrite layout.

Config.	R (mΩ)	L (μH)	Q (–)	M (μH)	k (–)	$\eta_{\max}$ (%)	Ferrite count		Cost (€)	$\Delta\eta/\text{Cost}$ ($\times 10^6$ % per €)
							#100	#150		
Air-core	70.13	132.02	1005.6	35.63	0.2699	99.24	–	–	0.00	–
Type I	121.24	213.81	941.9	68.00	0.3180	99.33	72	0	987.84	91.1
Type II	130.26	197.25	808.8	61.75	0.3130	99.21	64	0	878.08	–34.2
Type III	121.12	213.94	943.3	68.00	0.3178	99.33	0	48	1162.08	77.4
Type IV	127.79	217.05	908.0	75.80	0.3465	99.34	8	40	1078.16	92.8
Type V	121.59	198.30	871.6	60.90	0.3071	99.26	8	32	884.48	22.6
Type VI	114.46	146.94	685.6	31.11	0.2117	98.63	0	16	387.36	–157.5
Type VII	120.23	155.16	689.3	34.72	0.2238	98.71	8	16	497.12	–106.6
Type VIII	124.68	172.50	738.9	44.51	0.2580	98.96	0	28	677.88	–41.3

B.3 Coupling B

Table B.3: Coupling B: absolute electrical values and cost for each ferrite layout.

Config.	Tx R (mΩ)	Tx L (μH)	Tx Q (–)	Rx R (mΩ)	Rx L (μH)	Rx Q (–)	M (μH)	k (–)	$\eta_{\max}$ (%)	Ferrite count		Cost (€)	$\Delta\eta/\text{Cost}$ ($\times 10^6$ % per €)
										#100	#150		
Air-core	70.13	132.02	1005.6	32.44	44.07	725.5	15.57	0.2040	98.82	–	–	0.00	–
Type I	120.08	210.86	937.8	49.00	68.55	747.1	25.48	0.2120	98.88	29	20	882.04	68.0
Type II	120.08	210.86	937.8	49.05	66.45	723.5	25.52	0.2156	98.88	24	20	813.32	73.8
Type III	120.08	210.86	937.8	47.76	63.96	715.2	25.18	0.2169	98.88	20	20	758.60	79.1
Type IV	127.79	217.05	908.0	49.27	65.06	705.4	27.64	0.2393	98.89	8	28	787.64	88.9
Type V	120.08	210.86	937.8	45.41	48.71	572.9	17.45	0.1787	98.46	4	24	635.96	–56.6
Type VI	120.08	210.86	937.8	48.78	55.05	602.8	20.91	0.2003	98.66	8	24	690.84	–23.2

B.4 Coupling C

Table B.4: Coupling C: absolute electrical values and cost for each ferrite layout.

Config.	R (mΩ)	L (μH)	Q (–)	M (μH)	k (–)	$\eta_{\max}$ (%)	Ferrite count		Cost (€)	$\Delta\eta/\text{Cost}$ ($\times 10^6$ % per €)
							#100	#150		
Air-core	61.32	196.63	1712.7	52.67	0.2680	99.55	–	–	0.00	–
Type I	86.45	318.81	1969.5	107.00	0.3355	99.67	72	0	987.84	121.5
Type II	91.83	297.73	1731.6	92.63	0.3111	99.63	64	0	878.08	91.1
Type III	86.50	318.83	1968.6	107.00	0.3355	99.67	0	48	1162.08	103.3
Type IV	90.00	323.26	1918.1	111.18	0.3425	99.69	8	40	1078.16	129.9
Type V	88.16	296.78	1797.9	90.30	0.3043	99.63	8	32	884.48	90.5
Type VI	91.34	221.89	1297.5	46.53	0.2097	99.27	0	16	387.36	–72.3
Type VII	92.82	234.87	1351.5	51.93	0.2211	99.33	8	16	497.20	–44.3
Type VIII	93.04	258.39	1483.2	65.53	0.2536	99.47	0	28	677.88	–11.8

B.5 Coupling D

Table B.5: Coupling D: absolute electrical values and cost for each ferrite layout.

Config.	Tx R (mΩ)	Tx L (μH)	Tx Q (–)	Rx R (mΩ)	Rx L (μH)	Rx Q (–)	M (μH)	k (–)	$\eta_{\max}$ (%)	Ferrite count		Cost (€)	$\Delta\eta/\text{Cost}$ ($\times 10^6$ % per €)
										#100	#150		
Air-core	61.32	196.63	1712.7	54.46	142.05	1638.9	33.97	0.2032	99.34	–	–	0.00	–
Type I	86.25	314.21	1954.7	78.71	223.04	1513.3	66.00	0.2493	99.43	29	20	882.04	102.0
Type II	86.25	314.21	1954.7	79.17	215.97	1456.9	61.50	0.2360	99.43	24	20	813.32	110.7
Type III	86.25	314.21	1954.7	78.22	208.18	1421.4	55.80	0.2181	99.43	20	20	758.60	118.6
Type IV	90.00	323.26	1918.1	80.04	211.70	1412.7	59.97	0.2361	99.44	8	28	787.64	127.0
Type V	86.25	314.21	1954.7	79.08	157.80	1065.7	37.55	0.1750	99.18	4	24	635.96	–25.2
Type VI	86.25	314.21	1954.7	82.10	178.67	1162.3	45.19	0.1970	99.30	8	24	690.84	–5.8