



Università degli Studi di Padova

DEPARTMENT OF INFORMATION ENGINEERING

MASTER THESIS IN ICT FOR INTERNET AND MULTIMEDIA

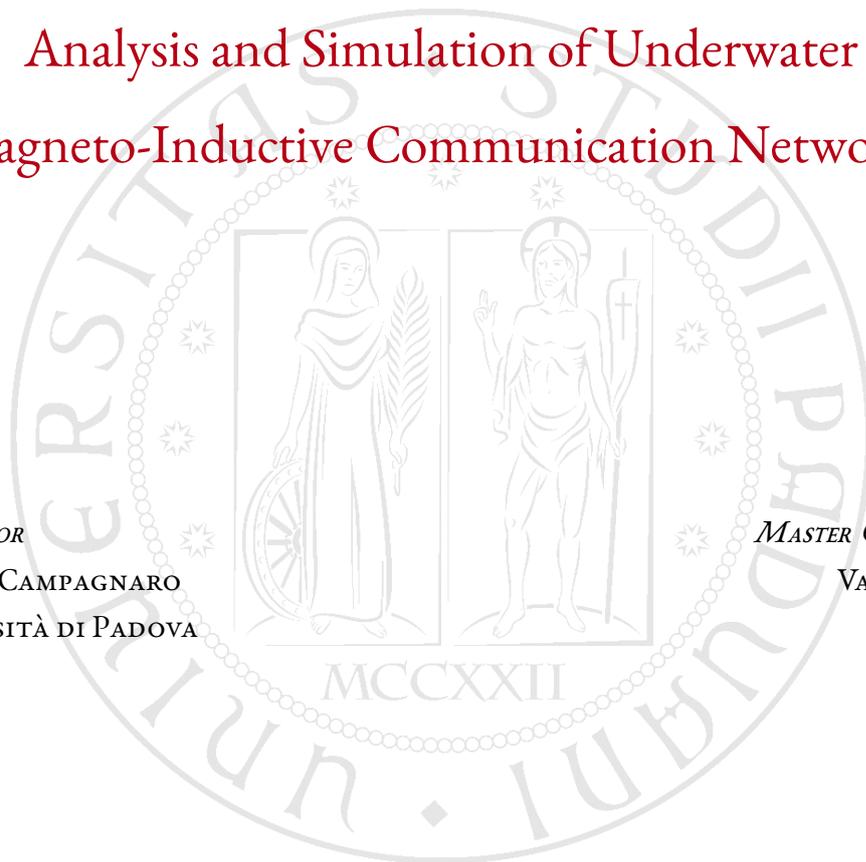
Analysis and Simulation of Underwater Magneto-Inductive Communication Networks

SUPERVISOR

FILIPPO CAMPAGNARO
UNIVERSITÀ DI PADOVA

MASTER CANDIDATE

VANESA ZIU



To my family, for you I walked this whole road.

To my mother, who placed in me a dream and a purpose no one can ever take away.

To my sister Ana, my hand-in-hand companion in this life, the blessing I count twice.

To Bora, for whom every effort is made, so she may know that it can be done.

To the generations before me, who carried weight so that I might carry books.

To Jesus, source of everything good in me.
May this work be my quiet thank you.

Abstract

This thesis develops and validates a physically grounded, simulation-ready model of underwater magneto-inductive (MI) communication and embeds it in the DESERT Underwater framework. Underwater communication is a key enabler for long-term environmental monitoring, infrastructure inspection and support to autonomous vehicles, but acoustic and optical technologies suffer from fundamental limitations in terms of capacity, latency and robustness. MI links, based on the quasi-static magnetic field generated by resonant coils, offer a channel that is almost free of multipath and fast fading, but their range is strongly limited by the steep decay of the magnetic near field and by conductive losses in water.

The analytical model starts from Maxwell's equations in a conductive medium and uses the skin depth as the main length scale for field attenuation in water. Circular coils and simple axial stacks are approximated as magnetic dipoles, and their mutual inductance is combined with a two-layer description that separates fully submerged underwater-to-underwater links from links that cross the water-air boundary. This leads to an equivalent RLC link budget that relates coil geometry and water conductivity to received power and to the signal-to-noise ratio for uncoded binary phase-shift keying over an additive white Gaussian noise channel, through expressions that are simple enough to be evaluated inside a packet-level simulator.

The model is instantiated in DESERT through two new modules: a coupling-based MI propagation module and a dedicated physical layer (implemented in DESERT as `UwMiCouplingPropagation` and `UwMiPhy`), which together convert geometry and water conductivity into analytical estimates of bit and packet error rates while remaining compatible with the existing protocol stack. A Tcl scenario and an automated sweep and post-processing pipeline produce packet delivery ratio curves, signal-to-noise ratio and received power trends, range tables and rate-range frontiers with full traceability to individual simulation runs.

For a simple low-power MI modem with a 12 cm radius coil (40 turns), transmit power 0 dBm at 200 kHz, bit rate 1 kbps and a noise figure of about 18 dB, the framework predicts ranges of about 12 m for fully submerged links and 15 m for links that cross the surface in fresh water, and about 3.5 m and 4.5 m in sea water, for target packet delivery ratios between 0.8 and 0.9. Coaxial stacks of up to three coils per node extend coverage by roughly 30–40% even under a fixed copper-mass constraint, while sweeps over frequency, quality factor, noise figure and useful bandwidth confirm that geometry and water conductivity dominate the MI link budget and that front-end parameters play a secondary role. Overall, the thesis provides a compact bridge between near-field MI physics and packet-level simulation, delivering quantitative design guidelines for underwater MI modems and a reusable building block for future studies on coding, interference and multi-hop MI networks.

Sommario

QUESTA TESI SVILUPPA E VALIDA un modello fisicamente fondato della comunicazione magneto-induttiva (MI) in ambiente subacqueo e lo integra nel framework di simulazione DESERT Underwater. La comunicazione subacquea è essenziale per il monitoraggio ambientale, l'ispezione di infrastrutture e il supporto a veicoli autonomi, ma le tecnologie acustiche e ottiche soffrono di limiti strutturali in termini di capacità, latenza e robustezza. La comunicazione MI, basata sul campo magnetico quasi-statico generato da bobine risonanti, offre un canale praticamente privo di multipath e fading veloce, ma con forti vincoli di portata dovuti al rapido decadimento del campo di near-field e alle perdite conduttive dell'acqua.

Il modello analitico parte dalle equazioni di Maxwell in un mezzo conduttivo e utilizza la profondità di penetrazione (skin depth) come principale parametro di scala per l'attenuazione del campo in acqua. Bobine circolari e semplici stack assiali sono approssimati come dipoli magnetici, e la loro induttanza mutua è combinata con una descrizione a due strati che separa i collegamenti completamente sommersi da quelli che attraversano l'interfaccia acqua-aria. In questo modo si ottiene un budget di collegamento RLC equivalente che mette in relazione geometria delle bobine, conducibilità dell'acqua e potenza ricevuta con il rapporto segnale-rumore di una modulazione binaria non codificata su canale gaussiano bianco, mediante espressioni sufficientemente compatte da essere valutate all'interno di un simulatore a livello di pacchetto.

Il modello è istanziato in DESERT tramite due nuovi moduli: un modulo di propagazione MI basato sull'accoppiamento e un livello fisico dedicato (implementati in DESERT come UwMiCouplingPropagation e UwMiPhy), che insieme trasformano geometria e conducibilità dell'acqua in stime analitiche dei tassi di errore di bit e di pacchetto mantenendo la piena compatibilità con lo stack esistente. Uno scenario Tcl e una pipeline automatizzata di sweep e post-processing producono curve di rapporto di consegna dei pacchetti, andamenti del rapporto segnale-rumore e della potenza ricevuta, tabelle di portata e frontiere di compromesso tra distanza e bit rate, con tracciabilità completa alle singole simulazioni.

Per un modem MI semplice e a bassa potenza, con una bobina di raggio 12 cm (40 spire), potenza trasmessa pari a 0 dBm a 200 kHz, bit rate di 1 kbps e noise figure di circa 18 dB, il framework prevede portate dell'ordine di 12 m per collegamenti completamente sommersi e 15 m per collegamenti che attraversano la superficie in acqua dolce, e circa 3,5 m e 4,5 m in acqua salata, per valori obiettivo del rapporto di consegna dei pacchetti compresi tra 0.8 e 0.9. Stack coassiali fino a tre bobine per nodo estendono la portata di circa il 30-40% anche sotto vincolo di massa di rame costante, mentre sweep su frequenza, fattore di qualità, noise figure e banda utile confermano che geometria e conducibilità dominano il budget MI e che

i parametri di front-end hanno un ruolo secondario. Nel complesso, la tesi costruisce un ponte compatto e interpretabile tra la fisica MI di near-field e la simulazione a livello di pacchetto, offrendo linee guida quantitative per la progettazione di modem MI subacquei e un blocco costitutivo riutilizzabile per studi futuri su codifica, interferenza e reti MI multi-hop.

Contents

ABSTRACT	v
LIST OF FIGURES	xii
LIST OF TABLES	xv
LIST OF ABBREVIATIONS	i
1 INTRODUCTION	3
1.1 Problem statement and motivation	4
1.2 Objectives and research questions	5
1.3 Methodology and scope	6
1.4 Contributions	7
1.5 Thesis organisation	8
2 UNDERWATER MAGNETO-INDUCTIVE CHANNEL	11
2.1 Operating regime of underwater MI links	12
2.2 Electromagnetic propagation in conductive media	12
2.2.1 Complex propagation constant and skin depth	12
2.2.2 Two-layer water-air geometry	13
2.3 Near-field MI coupling between coils	15
2.3.1 Magnetic dipole approximation	15
2.3.2 Mutual inductance of simple coil arrays	15
2.3.3 Geometric near-field path loss	17
2.4 Equivalent circuit and link budget	17
2.4.1 Lumped RLC model of an MI link	17
2.4.2 Incorporating conductive loss	18
2.4.3 Overall link gain and received power	19
2.4.4 Quality factor and effective bandwidth	19
2.4.5 From received power to bit energy-to-noise ratio	20
2.4.6 BPSK error probability and packet error rate	21
2.5 Relation to existing MI channel models	21
2.6 Modelling assumptions, limitations and specific contributions	22
3 MI SYSTEM MODEL AND DESERT IMPLEMENTATION	25

3.1	System model and assumptions	26
3.1.1	Node geometry and media	26
3.1.2	Coil and array parameters	27
3.1.3	Noise model and baseband signalling	28
3.2	DESERT Underwater architecture	29
3.2.1	End-to-end stack composition	29
3.2.2	Spectral mask and carrier	29
3.2.3	Configuration via Tcl	31
3.3	MI propagation model: <code>UwMiCouplingPropagation</code>	32
3.3.1	Position handling and two-layer splitting	32
3.3.2	Mutual inductance of arrayed coils	32
3.3.3	Conductive loss and skin depth	33
3.3.4	Overall link gain	33
3.4	MI physical layer: <code>UwMiPhy</code>	34
3.4.1	Transmit path	34
3.4.2	Receive path and acquisition test	35
3.4.3	Effective bandwidth and thermal noise	36
3.4.4	Mapping SNR to PER	36
3.5	Simulation scenarios and logging	37
3.5.1	Default configuration	37
3.5.2	One-way link and watchdog	38
3.5.3	Run-time logging and CSV summaries	38
3.6	Experiment automation and parameter sweeps	39
3.7	Post-processing, confidence intervals and range extraction	40
3.7.1	Aggregation of PDR and Wilson intervals	41
3.7.2	Extraction of PHY metrics from logs	41
3.7.3	Theory-versus-simulation PER curves and design tables	42
3.8	Chapter summary	43
4	RESULTS	45
4.1	PHY validation in AWGN	46
4.2	Single-coil link performance	47
4.2.1	UW→UW and UW→AW links	47
4.2.2	Range at PDR thresholds	49
4.3	Impact of depth and water-air interface	50
4.4	Array gain with multi-coil antennas	52
4.4.1	PDR vs distance	53
4.4.2	Range vs number of coils	54
4.5	Sensitivity to frequency and hardware parameters	55
4.5.1	Range vs carrier frequency	55
4.5.2	Range vs coil quality factor	56

4.5.3	Range vs receiver noise figure	56
4.5.4	Rate–range frontier over bandwidth and bit rate	57
4.6	Discussion and chapter summary	58
5	CONCLUSIONS AND FUTURE WORK	65
5.1	Conclusions	65
5.2	Future work	68
A	MI SIMULATION FRAMEWORK AND REPRODUCIBILITY	71
A.1	Code components and roles	71
A.2	Running a single MI simulation	72
A.3	Automated sweeps with <code>sweep_runs_v2.sh</code>	73
A.4	Post–processing and figure generation	74
A.5	Adapting the framework to other MI designs	75
	REFERENCES	76
	ACKNOWLEDGMENTS	79

Listing of figures

2.1	Node geometry and water–air interface for underwater UW→UW and UW→AW MI links.	14
2.2	Single–coil MI link and coaxial coil stacks per node. The total mutual inductance $M_{\text{sum}}(D)$ is obtained by summing the pairwise mutual inductances $M(d_{ij})$ between all transmit and receive coils.	16
2.3	Equivalent lumped RLC model of the magneto–inductive link: the transmit and receive resonant loops, characterised by (R_t, L_t, C_t) and (R_r, L_r, C_r, R_L) , are coupled by the mutual inductance M	18
3.1	End–to–end DESERT stack used in the simulations. Node 1 (Tx) runs the standard CBR–UDP–IP–CSMA/ALOHA stack downward towards the MI PHY, while Node 2 (Rx) processes packets upward from the MI PHY to the sink application. The custom MI modules (UwMiPhy and UwMiCouplingPropagation, shaded) interface with the common electromagnetic channel.	30
3.2	Automation pipeline for MI simulations. The shell script <code>sweep_runs_v2.sh</code> sweeps the design space and repeatedly calls the Tcl scenario <code>test_uwmi_net.tcl</code> , which configures DESERT/ns-2 runs with UwMiPhy and UwMiCouplingPropagation. Each run produces CSV summaries and per–packet UWMI_METRIC logs, which are parsed by the Python script <code>make_core_plots.py</code> to generate PDR curves, E_b/N_0 metrics and range tables used in Chapter 4.	44
4.1	PER as a function of E_b/N_0 for uncoded BPSK: analytical curve and DESERT simulations using UwMiPhy. Each marker aggregates multiple seeds at the same E_b/N_0 ; PER is plotted on a logarithmic scale.	47
4.2	PDR vs distance for the UW→UW link in freshwater ($\sigma = 0.05$ S/m) and seawater ($\sigma = 4.0$ S/m). Depths: $z_1 = z_2 = -5$ m.	48
4.3	PDR vs distance for the UW→AW link in freshwater and seawater. Depths: $z_1 = -2.0$ m, $z_2 = +1.0$ m.	49
4.4	Median E_b/N_0 vs distance for the UW→UW configuration at $\sigma = 0.05$ S/m and $\sigma = 4.0$ S/m. The dashed line indicates a conservative target $E_b/N_0 \approx 10$ dB for reliable BPSK reception.	50
4.5	Median E_b/N_0 vs distance for the UW→AW configuration.	51

4.6	Median received power as a function of distance for the UW→UW configuration. The dotted line shows an ideal $1/d^6$ near-field MI decay, anchored at the first point. The horizontal dashed line indicates the thermal-noise floor computed from the noise figure and effective bandwidth.	52
4.7	Median received power vs distance for the UW→AW configuration.	53
4.8	PDR vs distance for UW→AW in seawater, for different receiver heights above the surface z_2 . Conductivity $\sigma = 4.0$ S/m, $z_1 = -2$ m.	54
4.9	PDR vs distance for UW→AW in freshwater, for different z_2	55
4.10	Heatmap of PDR as a function of distance D and receiver height z_2 for the UW→AW link in seawater.	56
4.11	Heatmap of PDR as a function of distance D and receiver height z_2 for the UW→AW link in freshwater.	57
4.12	Array gain for UW→UW in freshwater ($\sigma = 0.05$ S/m): PDR vs distance for different numbers of coils per node in the “raw coil count” configuration (<code>auto_scale_R=0</code>).	58
4.13	Array gain for UW→UW in freshwater: PDR vs distance for different numbers of coils per node with equal total copper mass (<code>auto_scale_R=1</code> , effective series resistance scaled with N).	59
4.14	Range at $\text{PDR} \geq 0.8$ vs number of coils per node for the “raw coil count” configuration (UW→UW, $\sigma = 0.05$ S/m).	60
4.15	Range at $\text{PDR} \geq 0.8$ vs number of coils per node for the equal-copper-mass configuration.	61
4.16	Range at $\text{PDR} \geq 0.8$ vs carrier frequency f_0 for the UW→UW link in freshwater and seawater.	61
4.17	Range at $\text{PDR} \geq 0.8$ vs f_0 for the UW→AW link in freshwater and seawater.	62

Listing of tables

4.1	Range D_τ (in metres) for different PDR thresholds τ , link geometries and conductivities. Values are obtained by interpolation on the PDR vs distance curves.	50
4.2	Range $D_{0.8}$ vs quality factor Q for the UW \rightarrow UW link in freshwater ($\text{PDR} \geq 0.8$, $f_0 = 200$ kHz, $R_b = 1$ kbps).	62
4.3	Range $D_{0.8}$ vs receiver noise figure (NF) for the UW \rightarrow UW link in freshwater ($\text{PDR} \geq 0.8$, $f_0 = 200$ kHz, $R_b = 1$ kbps).	62
4.4	Range $D_{0.8}$ (in metres) for different (B, R_b) pairs in freshwater UW \rightarrow UW operation ($f_0 = 200$ kHz, $Q = 50$, $P_t = 0$ dBm).	63

List of Abbreviations

AWGN	Additive white Gaussian noise
BER	Bit error rate
BPSK	Binary phase–shift keying
CBR	Constant bit rate
CSMA	Carrier sense multiple access
CSV	Comma–separated values
DESERT	DEsign, Simulate, Emulate and Realize Testbeds (underwater networking framework)
E_b/N_0	Bit energy–to–noise power spectral density ratio
IP	Internet Protocol
MAC	Medium access control
MI	Magneto–inductive
MLL	Multi–link layer
NF	Noise figure
ns-2	Network Simulator 2
PDR	Packet delivery ratio
PER	Packet error rate
PHY	Physical layer
RF	Radio–frequency
RLC	Resistor–inductor–capacitor
Rx	Receiver
SNR	Signal–to–noise ratio
Tx	Transmitter
UDP	User Datagram Protocol
UW→AW	Underwater–to–air (MI link type)
UW→UW	Underwater–to–underwater (MI link type)

“Between Maxwell’s equations and underwater networks there is a long chain of models, approximations and implementations. This thesis walks that chain for magneto–inductive links.”

1

Introduction

Reliable underwater communication is a key enabler for many emerging applications, from long–term environmental monitoring and infrastructure inspection to autonomous vehicles and diver assistance in underwater sensor networks and marine monitoring systems [1, 2, 3, 4]. Despite decades of research, however, underwater links remain notoriously difficult to design. Acoustic systems achieve long ranges but suffer from low data rates, large propagation delays and strong multipath; optical links can offer high throughput but require very clear water, careful pointing and short distances; low–frequency radio waves are heavily attenuated by conductive seawater and rapidly become impractical beyond a few metres. There is therefore a persistent need for communication methods that can support short– to medium–range links with moderate data rates, low latency and robust performance in realistic water conditions.

Magneto–inductive (MI) communication is one such candidate [5, 6, 7]. Instead of radiating propagating waves, MI systems exploit the quasi–static magnetic field generated by resonant coils operating in the near field. At frequencies of a few tens to a few hundreds of kilohertz and distances of a few metres up to a few tens of metres, the magnetic field is dominated by a $1/d^3$ decay, leading to received power that scales approximately as $1/d^6$ with distance d . This behaviour is characteristic of near–field MI links based on small loop antennas [8, 9]. This steep near–field attenuation is a severe constraint on range, but it also yields a channel that is largely free of multipath and fast fading, and that is relatively insensitive to boundaries and small perturbations in the environment. Moreover, magnetic fields are

weakly affected by most natural underwater obstacles that would shadow acoustic or optical links.

The price of this robustness is that geometry and conductivity completely dominate the link budget. In an MI system, coil size, number of turns, orientation, and node separation determine the mutual inductance; water conductivity introduces frequency-dependent conductive loss via the skin depth. Electronics—amplifiers, matching networks, front-end noise—can only operate within the constraints imposed by this physical layer. Designing an underwater MI modem therefore requires a clear, quantitative understanding of how geometry and medium properties map into received power, the bit energy-to-noise power spectral density ratio E_b/N_0 , and packet error rate (PER) at the system level [8, 9, 5, 7].

1.1 PROBLEM STATEMENT AND MOTIVATION

The starting point of this thesis is the observation that, while the physics of MI coupling is well understood at the level of near-field electromagnetics and circuit models, there is still a gap between this physical description and packet-level performance evaluation in realistic underwater networking stacks.

On one side, analytical MI models typically focus on mutual inductance, skin depth and resonant coupling, providing closed-form expressions for path loss and received power under controlled assumptions [8, 9, 5, 7]. On the other side, underwater networking simulators such as DESERT Underwater [10] (henceforth DESERT) offer rich protocol stacks and channel abstractions, but they do not natively include MI-specific propagation models or physical layers. As a result, system designers who wish to evaluate MI networks in a realistic protocol environment are often forced to rely on rough approximations, ad-hoc link budgets or simplified additive white Gaussian noise (AWGN) channels decoupled from the underlying coil geometry and medium. At the same time, existing MI hardware work has mainly concentrated on coil and transceiver design and on medium-access protocols for small testbeds, rather than on simulation-ready models integrated into underwater network stacks [11, 12, 13].

This gap between physical models and network-level tools leads to the following practical question, which guides the rest of the thesis:

For a simple, low-power underwater MI modem based on circular coils, what link distances and packet delivery ratios are achievable in fresh and salt water, for both underwater-to-underwater and underwater-to-air geometries, once the modem is embedded in a realistic networking stack?

Answering this question requires bridging three levels of description:

1. a *physical* model of MI coupling and conductive loss, grounded in Maxwell's equations and lumped resistor–inductor–capacitor (RLC) circuits;
2. a *system–level* implementation that maps this model into a channel gain and noise figure inside a packet–level simulator;
3. a *numerical* study that explores how frequency, conductivity, coil arrays and front–end parameters shape the achievable range and reliability.

The present work focuses on constructing this bridge in a minimal yet physically faithful way, so that the resulting framework can support both performance evaluation and early–stage design decisions for underwater MI links.

1.2 OBJECTIVES AND RESEARCH QUESTIONS

The overall objective of the thesis is to develop and validate a physically grounded, simulation–ready model of underwater magneto–inductive communication, and to use it to quantify the design space of short–range MI links in realistic water conditions.

In more operational terms, the thesis is organised around the following concrete objectives:

- MI channel modelling: derive a simple analytical description of the underwater MI channel that starts from near–field coupling between coils and includes conductive loss via the skin depth, yielding expressions for link gain, received power and E_b/N_0 that can be plugged directly into a packet–level simulator.
- System integration in DESERT: embed this model in the DESERT Underwater framework by developing a dedicated MI propagation module and a custom MI physical (PHY) layer that exposes analytical binary phase–shift keying (BPSK) performance through the standard DESERT PHY interface.
- Numerical evaluation and design insight: perform a systematic set of simulations to characterise the achievable range of simple MI modems in representative freshwater and seawater scenarios, and to quantify the impact of key design parameters such as node depth, water–air interface, coil arrays, quality factor, noise figure and occupied bandwidth.

These objectives can be summarised in a set of guiding research questions:

- For a single circular coil with realistic dimensions and transmit power, what ranges are achievable at target packet delivery ratios in fresh and salt water?
- How does the presence of a water–air interface, as in underwater–to–air (UW→AW) links, change the MI link budget compared to fully submerged underwater–to–underwater (UW→UW) links?
- To what extent can per–node multi–coil arrays increase the coverage of MI links, both when extra copper is allowed and when the total copper mass is kept approximately fixed?
- How sensitive is the range to hardware parameters such as carrier frequency, coil quality factor and receiver noise figure, compared to the dominant effects of geometry and conductivity?
- How are bit rate and effective bandwidth traded against distance in MI links, and what rate–range relationship emerges for simple, low–power modem designs?

Addressing these questions requires careful alignment between the physical model, the simulator implementation and the post–processing of numerical results; a central theme of the thesis is to maintain this alignment explicitly and transparently.

1.3 METHODOLOGY AND SCOPE

The approach taken in this thesis is deliberately structured and minimalistic. Rather than aiming for the most general full–wave solution of Maxwell’s equations in complex underwater environments, the work focuses on a regime that is both practically relevant and mathematically tractable:

- *Near–field operation:* coils are small compared to their separation and the wavelength in the medium, so that the magnetic field can be approximated by a dipole model with $1/d^3$ decay and the received power scales as $1/d^6$.
- *Homogeneous media with a flat interface:* water is modelled as a homogeneous conducting medium with constant conductivity, and when a water–air interface is present it is taken to be flat and infinite. A two–layer approximation is used to split the transmitter–receiver segment into submerged and above–water portions.
- *Lumped RLC representation:* each node is modelled as a series RLC circuit at the operating frequency, with mutual inductance providing the coupling between coils and effective resistances collecting matching effects and practical implementation losses.

- *Analytical BPSK over AWGN PHY*: the physical layer is treated as uncoded BPSK over an equivalent AWGN channel, so MI propagation only changes the received power and hence E_b/N_0 ; the bit error rate (BER) and packet error rate (PER) are then obtained from standard textbook formulas.

Within this scope, the methodology is as follows.

First, the thesis develops a self-contained MI channel model that starts from Maxwell's equations in conductive media, introduces the skin depth as a key scaling parameter, and derives the mutual inductance between circular coils in the near field. This is combined with an equivalent RLC link budget that maps geometry, coil parameters and conductivity into link gain, received power and E_b/N_0 .

Second, the analytical model is brought into DESERT Underwater through two dedicated modules. `UwMiCouplingPropagation` evaluates the MI link gain from coil parameters, node geometry and water conductivity, including the two-layer $UW \rightarrow UW$ / $UW \rightarrow AW$ description. `UwMiPhy` takes the resulting received power, converts it into E_b/N_0 , applies the analytical BPSK-over-AWGN error model and exposes the corresponding packet error probabilities to the rest of the DESERT stack. A Tcl script acts as a light configuration front end for geometry, hardware and simulation parameters.

Third, a small set of shell and Python scripts is used to automate large parameter sweeps and to post-process the resulting logs into statistically robust performance metrics. Packet-level simulations are aggregated across seeds using Wilson confidence intervals, and range is consistently defined as the largest distance at which the empirical packet delivery ratio (PDR) stays above a chosen threshold. The same logs are also used to reconstruct E_b/N_0 and received-power trends, allowing direct comparison between simulation and analytical predictions.

The scope of the thesis is correspondingly delimited. Effects such as lateral coil misalignment, time-varying geometries, complex seabed or surface roughness, coding schemes beyond uncoded BPSK, and multi-hop network protocols are intentionally left to future work. The goal here is to construct a clean, well-understood MI building block that can later be extended and embedded into larger underwater systems.

1.4 CONTRIBUTIONS

Within the methodology and scope outlined above, the main contributions of the thesis are:

- Simulation-ready MI channel model. The thesis develops a compact, closed-form model of the underwater MI channel that combines: (i) dipole-based mutual inductance between circular coils and simple coil arrays; (ii) a two-layer conductive loss formulation based on the underwater segment of the link and the skin depth; and (iii) an equivalent RLC link budget that yields received power and E_b/N_0 in a form directly usable by a packet-level simulator.
- DESERT implementation of MI propagation and PHY. The work introduces two new modules in the DESERT Underwater framework: `UwMiCouplingPropagation`, an MI-specific propagation class that returns link gain as a function of geometry, coil parameters and conductivity, and `UwMiPhy`, a custom BPSK physical layer that maps this gain into analytical BER and PER. These modules are integrated into a minimal end-to-end stack and exposed through Tcl configuration parameters.
- Automated simulation pipeline and post-processing. The thesis builds an end-to-end automation pipeline around DESERT that combines a shell script for parameter sweeps, a Tcl scenario for run-time configuration, and Python scripts for log parsing, aggregation and plotting. Together, these scripts generate PER-versus- E_b/N_0 validation curves, PDR-versus-distance grids, range tables and rate-range frontiers, with each plotted point traceable to the corresponding set of simulation runs.
- Quantitative characterisation of underwater MI links. Using this framework, the thesis quantifies the achievable range of simple, low-power MI modems in freshwater and seawater for both UW→UW and UW→AW links, including the impact of node depth and water-air interface. It also evaluates array gains for multi-coil stacks, and assesses the sensitivity of range to carrier frequency, quality factor, noise figure and effective bandwidth. The results highlight the dominant role of geometry and conductivity, and show that moderate variations in front-end parameters have only second-order effects on coverage in the considered regime.

Beyond their immediate numerical values, these contributions provide a physically interpretable and practically usable framework for reasoning about underwater MI links at the system level. They can support early design decisions for MI modems and networks, and they serve as a basis for more advanced studies that include, for example, coding, interference or multi-hop topologies.

1.5 THESIS ORGANISATION

The remainder of the thesis is organised as follows.

- Chapter 2 develops the physical description of the underwater magneto–inductive channel. It starts from electromagnetic propagation in conductive media, introduces the skin depth, models coils as magnetic dipoles, derives mutual inductance and near–field path loss for single coils and simple arrays, and then assembles these pieces into an equivalent RLC link budget that yields received power, E_b/N_0 and packet error rate. The main modelling assumptions and limitations are stated explicitly.
- Chapter 3 presents the system model and DESERT implementation. It introduces the custom `UwMiCouplingPropagation` and `UwMiPhy` modules, describes their integration into the DESERT stack, and details the Tcl configuration interface, simulation scenarios and automation scripts used to generate large ensembles of runs.
- Chapter 4 reports the numerical results obtained with the proposed framework. It first validates the MI PHY against analytical BPSK performance in AWGN, then characterises the range and reliability of single–coil links in freshwater and seawater, studies the effect of depth and water–air interface, quantifies array gains for multi–coil antennas, and explores the sensitivity of range to frequency, quality factor, noise figure and bandwidth, including rate–range trade–offs.
- Chapter 5 gathers the main findings, formulates practical design guidelines for underwater MI links, and sketches directions for further work, such as richer coil geometries, more detailed noise and interference models, and the integration of MI links into larger multi–modal underwater network architectures.

With this structure, the thesis moves from physical modelling to software implementation and then to numerical evaluation, keeping a consistent thread: making underwater magneto–inductive communication quantitative and testable at the system level.

“In underwater magneto–inductive links, geometry and conductivity set the rules; electronics can only play within them.”

2

Underwater Magneto–Inductive Channel

This chapter develops the physical model of the underwater magneto–inductive (MI) channel that underpins the system–level implementation of Chapter 3 and the numerical results of Chapter 4. The goal is to start from Maxwell’s equations and basic circuit theory and arrive at a compact set of expressions for path loss, received power and signal–to–noise ratio (SNR) that can be used directly in a packet–level simulator.

The exposition proceeds in four steps. Section 2.1 delineates the operating regime of underwater MI communication and contrasts it with acoustic and radiowave propagation. Section 2.2 recalls electromagnetic wave propagation in conductive media and introduces the skin depth (2.3) as a key scaling parameter for underwater links [14]. Section 2.3 models the coils as magnetic dipoles and derives the mutual inductance and near–field path loss, following standard MI channel models originally developed for underground and underwater scenarios [8, 9]. Section 2.4 combines these ingredients into an equivalent lumped–element link model that yields received power and E_b/N_0 , which is exactly the quantity used by the custom PHY in Chapter 3. Section 2.5 briefly relates this model to existing MI channel formulations. Finally, Section 2.6 summarises the main modelling assumptions and limitations.

Throughout the chapter, the emphasis is on simple, closed–form expressions that are faithful to the approximations implemented later in DESERT, rather than on the most general full–wave solution of Maxwell’s equations.

2.1 OPERATING REGIME OF UNDERWATER MI LINKS

Underwater communication can be realised acoustically (pressure waves), by low-frequency radiowaves, by optical links in clear water, or by exploiting the quasi-static magnetic field generated by inductive coils. Magneto-inductive systems typically operate at carrier frequencies in the range of a few tens to a few hundreds of kilohertz, with node separations on the order of a few metres up to a few tens of metres [8, 9, 5, 7].

In this regime two features are particularly important.

NEAR-FIELD OPERATION. The wavelength of an electromagnetic wave in seawater is very short. At $f_0 = 200$ kHz, the wavelength in a good conductor is on the order of a few metres, and the reactive near field of a small loop extends only over a fraction of that distance [14, 8, 9]. MI links of interest here therefore operate deep in the near field: the magnetic field is dominated by a quasi-static term that decays as $1/d^3$ with separation d , rather than by radiated $1/d$ waves. As a consequence the channel is largely free of multipath and fast fading, but the received power decays very steeply, approximately as $1/d^6$.

CONDUCTIVE LOSS. Water conductivity introduces frequency-dependent attenuation. In freshwater, with typical conductivity $\sigma \approx 0.05$ S/m, this loss is modest at short ranges. In seawater, with σ on the order of 4 S/m, induced conduction currents dissipate energy and limit the achievable range much more severely. These conductive effects are conveniently summarised by the *skin depth*, introduced next [14].

The combination of near-field coupling and conductive loss therefore produces a channel that is strongly constrained by geometry and frequency, and that interacts in a non-trivial way with coil design, quality factor and front-end noise. The purpose of the rest of this chapter is to express these dependencies in an explicit, simulation-ready form that can be mapped directly into the DESERT Underwater framework (Chapter 3).

2.2 ELECTROMAGNETIC PROPAGATION IN CONDUCTIVE MEDIA

2.2.1 COMPLEX PROPAGATION CONSTANT AND SKIN DEPTH

Consider an electromagnetic wave with angular frequency $\omega = 2\pi f$ propagating in a homogeneous, isotropic medium with permittivity ε , permeability μ and conductivity σ . Under harmonic excitation, Maxwell's equations lead to a wave equation with complex propagation

constant [14]

$$\gamma = \alpha + j\beta = \sqrt{j\omega\mu(\sigma + j\omega\varepsilon)}, \quad (2.1)$$

where α (in nepers per metre) is the attenuation constant and β is the phase constant. In good conductors at low radio frequencies—a regime that includes seawater in the 10–500 kHz band—one typically has $\sigma \gg \omega\varepsilon$, so that the propagation constant can be approximated as

$$\alpha \approx \beta \approx \sqrt{\frac{\pi f \mu \sigma}{2}}. \quad (2.2)$$

A convenient parameter in this regime is the *skin depth* δ , defined as the distance over which the field magnitude decays by a factor $1/e$:

$$\delta = \sqrt{\frac{2}{\omega\mu\sigma}}. \quad (2.3)$$

For a uniform plane wave in a good conductor, the field envelope decreases as $e^{-\alpha d} \approx e^{-d/\delta}$, so that the power carried forward decays approximately as $e^{-2d/\delta}$ with propagation distance d [14]. At $f_0 = 200$ kHz and $\mu \approx \mu_0$, this yields a skin depth on the order of 5 m for $\sigma = 0.05$ S/m (fresh water) and about 0.6 m for $\sigma = 4.0$ S/m (seawater). In freshwater the skin depth is therefore larger than the typical link distance; in seawater it can be comparable to or smaller than the separation between nodes.

MI links do not behave as plane waves, but the skin depth is still a handy scaling parameter: it sets the distance over which medium-induced currents dissipate energy and therefore tells us how much conductive loss changes the nominal $1/d^6$ geometric decay.

2.2.2 TWO-LAYER WATER-AIR GEOMETRY

Two basic geometries are considered throughout the thesis.

- In underwater-to-underwater (UW→UW) links, both transmitter and receiver are located below the surface, so that the region where the magnetic field is appreciable lies entirely inside the conducting medium.
- In underwater-to-air (UW→AW) links, one node is placed below the surface while the other is above it, so that the magnetic field extends through both water and air.

Figure 2.1 illustrates these two geometries.

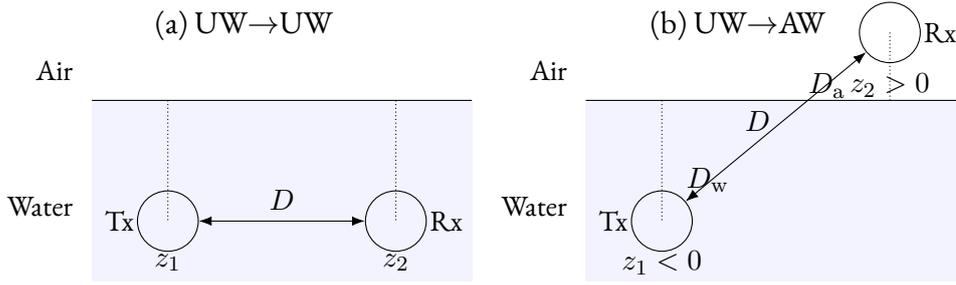


Figure 2.1: Node geometry and water–air interface for underwater UW→UW and UW→AW MI links.

A rigorous treatment of a water–air interface would require solving Maxwell’s equations with appropriate boundary conditions, accounting for reflections, refractions and lateral variations at the interface [14]. Such a treatment is unnecessarily complex for the short ranges and simple coil geometries considered here. Instead, a tractable *two-layer* model is adopted, consistent with the implementation of Chapter 3 and with simplified underwater MI interface treatments in the literature [9].

Let D be the centre–to–centre distance between the two coils. The straight line joining their centres is split into an *underwater* segment of length D_w and an *air* segment of length D_a , with $D_w + D_a = D$. Conductive attenuation is applied only along D_w according to the skin–depth model, whereas the portion in air is treated as essentially lossless over the short distances considered here.

This yields a simple path–loss factor that can be applied multiplicatively to the near–field coupling and that has a clear geometric interpretation: only the submerged segment contributes to conductive loss, whereas the air segment contributes only to the $1/D^3$ geometric decay.

One might wonder whether an additional loss term should be included at the water–air boundary itself, by analogy with acoustic or optical links where the interface is a strong reflector. In the quasi–static MI regime considered here this is not the case. The key parameter for the magnetic boundary conditions is the permeability μ , and both air and (non–magnetic) water have $\mu_r \approx 1$, so the normal component of \mathbf{B} and the tangential component of \mathbf{H} remain essentially continuous across the interface. As a result, the magnetic near field crosses the water–air boundary without a large discrete reflection; the main effect of the interface is that the submerged part of the path experiences conductive attenuation, while the segment in air does not. This is consistent with the MI literature, which explicitly notes that the communication quality of MI/RF links is “not disturbed” by the water/air boundary because

of the similar magnetic permeabilities of the two media [3, 9, 7]. Within the short ranges (a few to a few tens of metres) and frequencies (tens to a few hundreds of kilohertz) considered in this thesis, modelling the interface via the two-layer (D_w, D_a) split and applying the skin-depth factor only on D_w is therefore a reasonable approximation; more detailed layered-medium solutions are left to future work.

2.3 NEAR-FIELD MI COUPLING BETWEEN COILS

2.3.1 MAGNETIC DIPOLE APPROXIMATION

A circular current loop of radius a carrying a time-harmonic current $Ie^{j\omega t}$ can be modelled, at distances large compared with a , as a magnetic dipole with magnetic moment

$$\mathbf{m} = NIA \hat{\mathbf{n}}, \quad (2.4)$$

where N is the number of turns, $A = \pi a^2$ is the loop area and $\hat{\mathbf{n}}$ is a unit vector normal to the loop plane. In the magneto-quasi-static regime ($r \ll \lambda$ and $r \ll \delta$), the dominant terms of the dipole field decay as $1/r^3$ with distance r from the source [14, 8, 9]. For a single transmit (Tx) coil and a single receive (Rx) coil, both circular and coaxial, separated by a centre-to-centre distance d much larger than their radii, the mutual inductance can be approximated by the classical dipole-dipole expression [8, 9]

$$M(d) = \frac{\mu N_t N_r A_t A_r}{2\pi d^3}, \quad A_t = \pi a_t^2, \quad A_r = \pi a_r^2, \quad (2.5)$$

where (N_t, a_t) and (N_r, a_r) are the number of turns and radii of the transmit and receive coils, respectively, and μ is the permeability of the surrounding medium. More exact expressions in terms of elliptic integrals exist for arbitrary geometry, but (2.5) captures the essential $1/d^3$ distance dependence and is accurate when $d \gg a_t, a_r$ [8].

2.3.2 MUTUAL INDUCTANCE OF SIMPLE COIL ARRAYS

Many practical MI systems employ more than one coil per node to increase the magnetic moment or to shape the field. In this thesis, and in the associated DESERT implementation, a node-level array is modelled as $N_{t,\text{coils}}$ transmit coils and $N_{r,\text{coils}}$ receive coils, all circular, coaxial and aligned with the link axis.

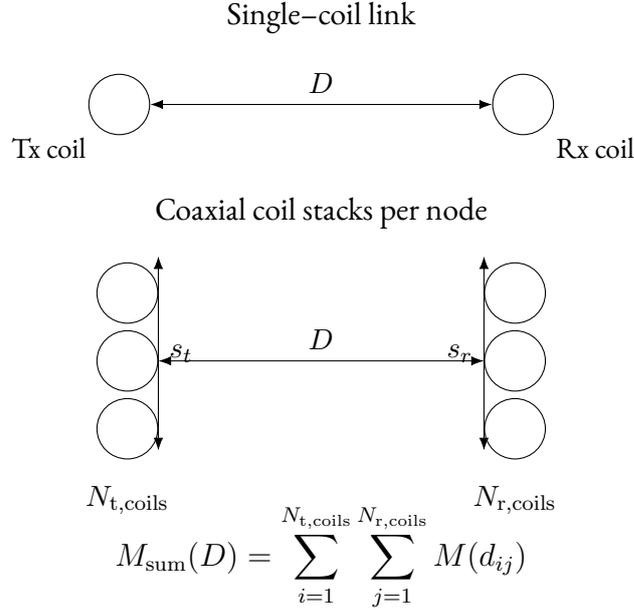


Figure 2.2: Single-coil MI link and coaxial coil stacks per node. The total mutual inductance $M_{\text{sum}}(D)$ is obtained by summing the pairwise mutual inductances $M(d_{ij})$ between all transmit and receive coils.

The centres of the coils belonging to the same node are placed at positions

$$z_i^{(t)} = \left(i - \frac{N_{\text{t,coils}} - 1}{2} \right) s_t, \quad i = 0, \dots, N_{\text{t,coils}} - 1, \quad (2.6)$$

and analogously for the receiver, with spacing s_r . For a given centre-to-centre separation D between the two nodes, the axial distance between the i -th transmit coil and the j -th receive coil is $d_{ij} = |D + z_j^{(r)} - z_i^{(t)}|$. Under the same dipole approximation, the total mutual inductance seen by the link is taken as the sum of all pairwise contributions:

$$M_{\text{sum}}(D) = \sum_{i=1}^{N_{\text{t,coils}}} \sum_{j=1}^{N_{\text{r,coils}}} M(d_{ij}) = \sum_{i=1}^{N_{\text{t,coils}}} \sum_{j=1}^{N_{\text{r,coils}}} \frac{\mu N_t N_r A_t A_r}{2\pi d_{ij}^3}. \quad (2.7)$$

Intra-array mutual coupling (between coils on the same node) is not modelled explicitly; its effect is absorbed into effective series resistances and an aggregate coupling factor introduced in Section 2.4. This is consistent with many MI system models, where node-level arrays are treated as an “equivalent dipole” with an increased magnetic moment [8, 9].

2.3.3 GEOMETRIC NEAR-FIELD PATH LOSS

In the absence of conductive loss, the $1/d^3$ dependence of $M(d)$ in (2.5) translates into a received power that falls as $1/d^6$. As shown more explicitly in Section 2.4, for two matched RLC loops with mutual inductance $M_{\text{sum}}(D)$ the power transfer ratio contains a factor $|M_{\text{sum}}(D)|^2$, so that in the dipole approximation

$$P_r(D) \propto \frac{1}{D^6}. \quad (2.8)$$

This extremely steep decay is a distinctive feature of near-field MI links: doubling the separation reduces the received power by roughly $6 \cdot 10 \log_{10} 2 \approx 18$ dB. As soon as conductive effects become significant, an additional exponential term appears and the effective path loss beyond several skin depths becomes even more severe, especially in seawater [9].

2.4 EQUIVALENT CIRCUIT AND LINK BUDGET

2.4.1 LUMPED RLC MODEL OF AN MI LINK

A magneto-inductive link between two nodes can be modelled as two resonant RLC circuits coupled by mutual inductance M , as shown in Fig. 2.3 [8, 9, 5, 7].

In its simplest form, the transmit side consists of a series resistance R_t , loop inductance L_t and tuning capacitance C_t , driven by a voltage source. The receive side consists of an analogous series resistance R_r , loop inductance L_r , capacitance C_r and a load. The inductors are coupled by M , so that a current I_t in the transmit loop induces a voltage $j\omega M I_t$ in the receive loop.

Under matched conditions and near resonance, small-signal analysis yields a power transfer ratio of the form

$$G_c(f, D) = \kappa^2 \frac{\omega^2 M_{\text{sum}}^2(D)}{4R_{t,\text{eff}} R_{r,\text{eff}}}, \quad (2.9)$$

where $M_{\text{sum}}(D)$ is the total mutual inductance (cf. (2.5)–(2.7)), $R_{t,\text{eff}}$ and $R_{r,\text{eff}}$ are effective resistances that include matching networks and the receive load, and κ is a dimensionless coupling factor that aggregates deviations from the ideal transformer model (additional insertion losses, mismatch, non-idealities). The factor $1/4$ reflects the optimal power transfer condition in a loosely coupled transformer.

Equation (2.9) represents the *geometric coupling gain* between the coils. In the absence of conductive loss it would be the only term in the path loss. In practice, the conducting

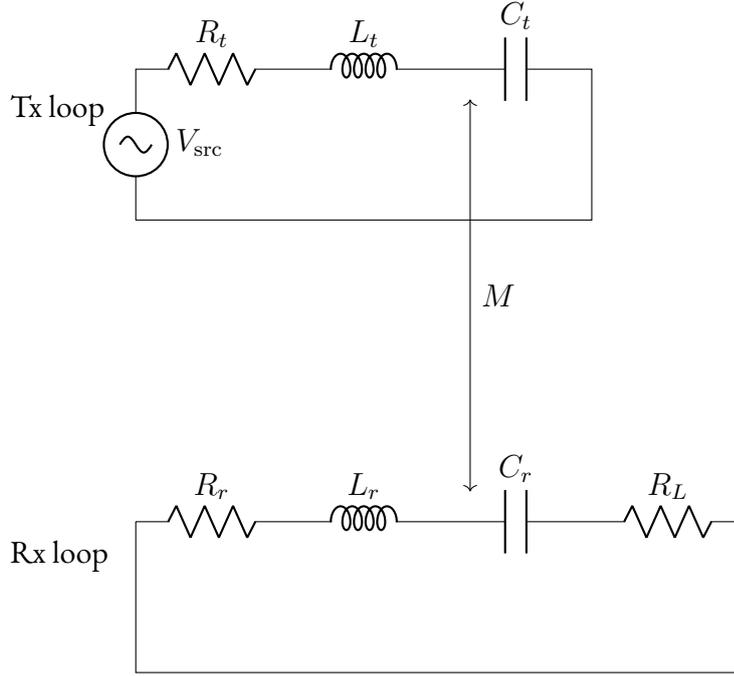


Figure 2.3: Equivalent lumped RLC model of the magneto-inductive link: the transmit and receive resonant loops, characterised by (R_t, L_t, C_t) and (R_r, L_r, C_r, R_L) , are coupled by the mutual inductance M .

medium introduces an additional factor associated with the skin depth.

2.4.2 INCORPORATING CONDUCTIVE LOSS

Following Section 2.2, conductive loss is modelled as an exponential attenuation over the submerged portion of the link. Let D_w denote the length of the segment between node centres that lies underwater. At frequency f and conductivity σ we approximate the medium-induced power transfer factor as

$$L_m(f, D_w) = \exp\left(-\frac{2D_w}{\delta(f, \sigma)}\right), \quad \delta(f, \sigma) = \sqrt{\frac{2}{\omega\mu\sigma}}, \quad (2.10)$$

where $\delta(f, \sigma)$ is the skin depth from (2.3) [14, 9]. For freshwater, with $\sigma = 0.05$ S/m, the exponent $2D_w/\delta$ remains small over the ranges considered and $L_m \approx 1$. For seawater, with $\sigma = 4.0$ S/m, the exponential term can substantially reduce the received power even when the geometric coupling is still relatively strong [9].

The two-layer geometry of Section 2.2 is then included simply by setting $D_w = D$ for

UW→UW links and $D_w < D$ for UW→AW links, so that only the submerged part of the path contributes to the exponential loss.

2.4.3 OVERALL LINK GAIN AND RECEIVED POWER

Combining geometric coupling and conductive loss, the overall link-level power gain is modelled as

$$G(f, D, D_w) = G_c(f, D) L_m(f, D_w), \quad (2.11)$$

with G_c given by (2.9) and L_m by (2.10). For a given transmit power P_t , the received power is therefore

$$P_r = G(f, D, D_w) P_t. \quad (2.12)$$

In the DESERT implementation of Chapter 3, this gain is clipped to a physically reasonable interval $[10^{-30}, 1]$ for numerical robustness; this clipping does not alter the analytical form of (2.11).

2.4.4 QUALITY FACTOR AND EFFECTIVE BANDWIDTH

Each resonant coil can be characterised by a quality factor

$$Q = \frac{\omega_0 L}{R}, \quad (2.13)$$

where $\omega_0 = 2\pi f_0$ is the resonance frequency of the RLC circuit, L is the loop inductance and R is the series resistance. The corresponding 3 dB bandwidth is approximately

$$B_{\text{res}} \approx \frac{f_0}{Q}. \quad (2.14)$$

A higher Q increases the peak transfer at f_0 but narrows the frequency range over which significant power is delivered [8, 9].

On the noise side, the front-end is assumed to be dominated by white thermal noise over the frequencies of interest. Noise outside the effective passband is filtered out by the coil resonance and subsequent stages. The *effective noise bandwidth* is therefore modelled as

$$B_{\text{eff}} = \min\left(B, \frac{f_0}{Q}\right), \quad (2.15)$$

where B is an explicit, configurable spectral mask bandwidth. Two limiting regimes are captured by (2.15). For low Q or narrow spectral masks, $B_{\text{eff}} \approx B$ and the front-end behaves as an ideal rectangular filter of width B . For high Q , the coil itself becomes the dominant band-pass element and $B_{\text{eff}} \approx f_0/Q$. This notion of effective bandwidth is used consistently in translating received power into SNR and matches the implementation of UWMiPhy in Chapter 3.

2.4.5 FROM RECEIVED POWER TO BIT ENERGY-TO-NOISE RATIO

The residual noise at the detector input is modelled as additive white Gaussian noise (AWGN) with double-sided power spectral density $N_0 = kTF$, where k is Boltzmann's constant, T is the equivalent noise temperature and $F = 10^{\text{NF}_{\text{dB}}/10}$ is the receiver noise factor associated with the receiver noise figure (NF) expressed in decibels as NF_{dB} [15].

The noise power over the effective bandwidth B_{eff} is

$$N = kTFB_{\text{eff}}. \quad (2.16)$$

Combining (2.12) and (2.16), the linear SNR at the detector input is

$$\text{SNR}_{\text{lin}} = \frac{P_r}{N} = \frac{G(f, D, D_w) P_t}{kTFB_{\text{eff}}}. \quad (2.17)$$

Let R_b denote the information bit rate. Assuming that the signal is confined to the same effective bandwidth B_{eff} over which noise is measured, the usual relation between SNR and bit-energy-to-noise-density ratio yields [15]

$$\left(\frac{E_b}{N_0} \right)_{\text{lin}} = \text{SNR}_{\text{lin}} \frac{B_{\text{eff}}}{R_b} = \frac{G(f, D, D_w) P_t}{kTFR_b}. \quad (2.18)$$

In other words, provided that B_{eff} is matched to the signal, E_b/N_0 depends on link gain, transmit power and bit rate, but not explicitly on the chosen front-end bandwidth.

This expression already anticipates several trends observed in Chapter 4: for fixed R_b and front-end noise parameters (T, F), the dominant design levers for range are the link gain $G(f, D, D_w)$, i.e., geometry, coil design and conductivity. Parameters such as quality factor Q , spectral mask bandwidth B and noise figure influence range only through second-order changes in G and F . This is reflected in the weak sensitivity of the simulated range to moderate variations in Q , NF and (B, R_b) .

In the DESERT implementation, an additional “noise margin” term can be included as a multiplicative factor on $(E_b/N_0)_{\text{lin}}$ to absorb unmodelled losses. For clarity this margin is not shown explicitly in (2.18) and is treated as an implementation detail in Chapter 3.

2.4.6 BPSK ERROR PROBABILITY AND PACKET ERROR RATE

The present work assumes uncoded binary phase–shift keying (BPSK) modulation. For coherent BPSK in AWGN, the bit error rate (BER) is given by the textbook expression [15]

$$P_b = \frac{1}{2} \operatorname{erfc} \left(\sqrt{\frac{E_b}{N_0}} \right), \quad (2.19)$$

where the argument is understood in linear units. For a packet of L bytes ($8L$ bits), and assuming independent bit errors, the corresponding packet error probability is

$$P_e = 1 - (1 - P_b)^{8L}. \quad (2.20)$$

Equations (2.11)–(2.18), together with (2.19) and (2.20), completely specify the analytical PER model used by the custom MI PHY: geometry, conductivity and coil design determine the link gain $G(f, D, D_w)$; front–end parameters map this gain into E_b/N_0 ; and the AWGN BPSK model provides P_e . Chapter 3 instantiates these expressions in the `UwMiCouplingPropagation` and `UwMiPhy` modules, and Chapter 4 validates the resulting PER curves against the analytical BPSK reference.

2.5 RELATION TO EXISTING MI CHANNEL MODELS

The structure of the model developed above is consistent with classical MI channel formulations for underground and underwater links, which also combine dipole–dipole mutual inductance, conductive loss and an equivalent RLC representation of the transceiver circuits [8, 9, 5, 7]. More general formulations may use exact mutual inductance expressions based on elliptic integrals, full three–dimensional coil layouts or full–wave numerical solvers for complex environments.

Within this broader family, the present work deliberately adopts the following simplifications: (i) the asymptotic $1/d^3$ mutual inductance of (2.5) instead of elliptic–integral expressions; (ii) a two–layer treatment of $UW \rightarrow UW$ and $UW \rightarrow AW$ links in terms of the underwater segment D_w ; and (iii) a one–dimensional axial array model based on the summed mutual

inductance $M_{\text{sum}}(D)$ in (2.7). These choices keep the expressions closed form and analytically transparent, and make them directly compatible with the DESERT Underwater framework, which did not previously include any MI-specific propagation model or PHY [10].

2.6 MODELLING ASSUMPTIONS, LIMITATIONS AND SPECIFIC CONTRIBUTIONS

The MI channel model used in this thesis is deliberately simple. It is designed to capture the dominant physics over the short ranges and frequencies of interest while remaining lightweight enough for large parameter sweeps in DESERT. The main assumptions are summarised below.

- Near-field regime and dipole approximation. Coils are assumed small compared to their separation and to the wavelength in the medium, so that the $1/d^3$ dipole-dipole expression (2.5) is adequate. Radiated far-field components, retardation effects and higher-order multipoles are neglected [8, 14, 9, 7].
- Coaxial, aligned coils. Transmit and receive coils are assumed coaxial and aligned with the link axis. Lateral misalignment, tilt and more complex array geometries are left to future work. Moderate misalignments would reduce the effective mutual inductance but would not change the qualitative $1/d^6$ decay.
- Homogeneous water column and flat interface. Water is modelled as a homogeneous conducting medium with constant conductivity σ and permeability μ . The water-air interface is assumed flat and infinite. Stratification, salinity gradients, surface roughness and seabed proximity are not considered [14].
- Two-layer conductive loss and interface. In UW→AW links, conductive loss is applied only to the part D_w of the link that lies in water, using (2.10), while the section in air is treated as lossless. No extra “interface loss” term is added at the water-air boundary. This is a deliberate choice: in the quasi-static MI regime the magnetic field is mainly governed by the permeability μ , which is essentially the same in air and in (non-magnetic) water, so the magnetic near field crosses the interface without a strong reflection penalty. The dominant loss mechanism is the exponential attenuation in the conducting layer, captured by D_w and the skin depth, rather than the boundary itself [9, 7]. A full stratified-medium solution with Sommerfeld integrals would be required for very long ranges or more complex geometries, and is left to future work.
- Lumped RLC representation. Each node is represented by a series RLC circuit at the operating frequency, with effective resistances $R_{t,\text{eff}}$ and $R_{r,\text{eff}}$ and a lumped coupling factor κ that absorb matching networks, load and implementation losses. Frequency

dependence of loop parameters beyond the resonance behaviour encapsulated in Q is not modelled explicitly [8, 9].

- Array simplifications. Multi-coil arrays are modelled as one-dimensional stacks of identical coils with fixed spacing along the link axis, all excited in phase with the same current. Intra-array coupling and three-dimensional layouts are captured only indirectly through effective resistances and κ , as in (2.7).
- AWGN noise model and uncoded BPSK. Receiver noise is modelled as white and Gaussian over the effective bandwidth B_{eff} . The baseband modulation is uncoded BPSK, with error probabilities given by (2.19) and (2.20) [15]. Other constellations, coding schemes and interference sources are not considered.

Within this simplified setting, the model adds three concrete elements to the thesis: (i) a two-layer UW→UW / UW→AW loss model expressed in terms of the underwater segment D_w ; (ii) a compact coil-array description based on the summed mutual inductance $M_{\text{sum}}(D)$ and effective resistances; and (iii) a direct mapping from this MI link budget to E_b/N_0 and packet error rate that matches the BPSK PHY interface used in DESERT. Together, these three pieces link the MI physics to the packet-level simulations presented in Chapters 3–4.

Despite its simplifications, the model reproduces the main trends reported in the MI literature for short-range underground and underwater links [8, 9, 5, 7]. In particular, it captures the $1/d^6$ near-field decay, the dominant role of conductivity via the skin depth, the influence of coil geometry and arrays on coupling gain, and the conversion from received power to E_b/N_0 and packet error rate embodied in (2.11)–(2.20).

The next chapter instantiates this physical model inside the DESERT Underwater framework by introducing a custom propagation module and an MI-specific PHY, and by building a simulation environment that maps the parameters introduced here (distance, conductivity, frequency, coil design and noise figure) into packet-level performance metrics.

“All models are wrong, but models that match both physics and packets are at least useful.”

Paraphrasing G. E. P. Box

3

MI System Model and DESERT Implementation

This chapter introduces the system model and software implementation used to simulate underwater magneto–inductive (MI) links in the DESERT Underwater framework [10].

The goal is to bridge the physical description of the MI channel, outlined in Chapter 2, with a practical packet–level simulator in which the impact of conductivity, coil geometry and hardware parameters can be studied under realistic networking stacks.

The proposed architecture is based on two custom modules:

- `UwMiCouplingPropagation`, a propagation class that computes the MI link gain from coil parameters, node geometry and water conductivity, including a two–layer underwater–to–air (UW→AW) model;
- `UwMiPhy`, a physical layer that maps the received power and noise figure into an equivalent E_b/N_0 and computes the packet error probability using the analytical BPSK over AWGN formula.

These modules are integrated into a minimal but complete application–to–PHY stack, configured by a Tcl script that exposes the main system parameters and writes results to comma–separated values (CSV) files for later analysis.

Beyond a straightforward configuration of existing DESERT components, this work contributes: (i) a new MI–specific propagation class with conductive two–layer support for both

UW→UW and UW→AW links; (ii) a custom MI physical layer that exposes analytical BPSK performance through a standard DESERT interface; and (iii) an automated sweep and post-processing pipeline that turns large ensembles of packet-level runs into statistically robust PDR, range and E_b/N_0 metrics. The performance of this architecture is evaluated in Chapter 4.

3.1 SYSTEM MODEL AND ASSUMPTIONS

We consider a single MI link between two DESERT nodes, each equipped with a circular coil antenna. Unless otherwise stated, the coils are assumed to be coaxial and aligned along the line joining the node centres, so that the magnetic coupling can be approximated by a dipole-dipole model in the near field.

The present implementation therefore focuses on coaxial, aligned coils in a homogeneous water column with a flat air-water interface. Effects such as lateral misalignment, tilt between coil axes, stratified water layers or complex boundary geometries are left to future work and are not explicitly modelled here.

3.1.1 NODE GEOMETRY AND MEDIA

Let the transmitter and receiver positions be $\mathbf{r}_1 = (x_1, y_1, z_1)$ and $\mathbf{r}_2 = (x_2, y_2, z_2)$ respectively, with $z = 0$ denoting the air-water interface (water below, air above). The Euclidean distance between the node centres is

$$D = \|\mathbf{r}_2 - \mathbf{r}_1\| = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2 - z_1)^2}. \quad (3.1)$$

Two basic geometries are used throughout the thesis:

- underwater-to-underwater (UW→UW), where both nodes are submerged, typically with $z_1 = z_2 = -5$ m;
- underwater-to-air (UW→AW), where the transmitter is submerged at $z_1 < 0$ and the receiver is above the surface at $z_2 > 0$.

In the latter case the segment between the nodes crosses the interface. The implementation therefore decomposes the straight line link into a water segment of length D_w and an air segment of length D_a , such that $D_w + D_a = D$. Conductive loss is applied only over D_w , while the magnetic field in air is treated as lossless over the short distances of interest.

In other words, at the implementation level the water–air boundary enters exactly through the computation of D_w : the medium is treated as piecewise homogeneous (water, then air), with conductive attenuation applied in water and negligible additional loss at the interface itself, consistently with the two–layer MI model of Chapter 2. The splitting logic is implemented in `UwMiCouplingPropagation::splitUnderwaterAir_()` and is enabled or disabled via a boolean flag `use_two_layer_`.

Throughout the thesis, two representative conductivities are used: $\sigma = 0.05$ S/m for fresh or clean water and $\sigma = 4.0$ S/m for salt water. The relative permeability is assumed to be $\mu_r \approx 1$, so that the absolute permeability is simply $\mu = \mu_0\mu_r$.

3.1.2 COIL AND ARRAY PARAMETERS

Each node is modelled as a stack of circular coils aligned with the link axis. For a given node, the per–coil parameters are:

- N_t, N_r : number of turns of the transmit and receive coils, respectively;
- a_t, a_r : coil radii;
- R_t, R_r : series resistances;
- κ : coupling factor collecting matching and radiation–efficiency effects;
- μ_r : relative permeability of the medium.

To capture the effect of multiple coils per node, the propagation module introduces the parameters `Nt_coils_` and `Nr_coils_`, denoted by $N_{t,\text{coils}}$ and $N_{r,\text{coils}}$ in formulas. The coils in a stack are spaced by s_t and s_r metres along the link axis (Tcl parameters `st` and `sr`), and their centres are symmetrically placed around the node centre. For a stack of N coils with spacing s , the axial offsets are

$$z_i = \left(i - \frac{N-1}{2}\right)s, \quad i = 0, \dots, N-1, \quad (3.2)$$

as implemented in `coilOffsets_()`.

Two design regimes are supported:

1. *Raw coil count (more copper)*: the effective series resistance of each node is kept equal to that of a single coil, i.e. $R_{t,\text{eff}} = R_t$ and $R_{r,\text{eff}} = R_r$. Adding coils thus increases the amount of copper and the magnetic moment without increasing ohmic loss.

2. *Equal copper mass*: the effective series resistances grow approximately in proportion to the number of coils, $R_{t,\text{eff}} \approx N_{t,\text{coils}}R_t$ and $R_{r,\text{eff}} \approx N_{r,\text{coils}}R_r$. This mimics a fixed copper budget that is shared across more, thinner coils.

The choice between these regimes is controlled by the flag `auto_scale_R_`: when it is set to 1, the resistance scaling is enabled; otherwise, R_t and R_r are left unchanged.

In the current implementation, all coils in a stack are identical and are driven in phase with the same current amplitude. Mutual coupling and possible pattern-shaping effects within the stack are not modelled explicitly; their combined impact is absorbed into the lumped coupling factor κ and the effective resistances $R_{t,\text{eff}}$ and $R_{r,\text{eff}}$. This keeps the coupling law close to the simple $1/d^3$ dipole approximation, while still letting us change the number of coils and their spacing to see how they affect the overall mutual inductance.

3.1.3 NOISE MODEL AND BASEBAND SIGNALLING

On the receive side, the thermal noise at the detector input is modelled as additive white Gaussian noise (AWGN) with double-sided power spectral density $N_0 = kTF$, where k is Boltzmann's constant, T is the equivalent noise temperature and F is the receiver noise factor. In practice, the effective noise power is computed as

$$N = kT_{\text{noise}}FB_{\text{eff}}, \quad (3.3)$$

where B_{eff} is the effective noise bandwidth of the overall front end.

The chosen baseband modulation is uncoded BPSK at a bit rate R_b . Rather than explicitly simulating waveforms, the PHY uses the analytical BPSK over AWGN expression for the bit error probability

$$P_b = \frac{1}{2} \operatorname{erfc}(\sqrt{E_b/N_0}), \quad (3.4)$$

where E_b/N_0 is the bit energy to noise spectral density ratio. This is the same expression used in the physical channel model of Chapter 2, now instantiated at the implementation level inside `UwMPhy`.

For a packet of L bytes ($8L$ bits), assuming independent bit errors, the packet error probability is

$$P_e = 1 - (1 - P_b)^{8L}. \quad (3.5)$$

This abstraction is implemented in `UwMPhy` and validated against analytical curves in Section 4.1.

3.2 DESERT UNDERWATER ARCHITECTURE

The implementation is based on the DESERT Underwater framework (DEsign, Simulate, Emulate and Realize Testbeds, DESERT), which extends Network Simulator 2 (ns-2) with a modular stack tailored to underwater networking.

Each node hosts a set of protocol modules connected in a layered fashion, with clear interfaces between application, transport, network, link and physical layer.

3.2.1 END-TO-END STACK COMPOSITION

The simulation script instantiates two nodes, each running the following stack:

- application: a constant bit rate (CBR) source (`Module/UW/CBR`);
- transport: User Datagram Protocol (UDP) (`Module/UW/UDP`);
- network: Internet Protocol (IP) (`Module/UW/IP`) with static routing;
- link: a simple medium access control (MAC) layer based on carrier-sense multiple access / ALOHA (CSMA/ALOHA) (`Module/UW/CSMA_ALOHA`);
- link-to-PHY glue: the multi-link layer (MLL; `Module/UW/MLL`);
- physical layer: `Module/UW/MI/PHY/Custom` (the custom `UwMiPhy` module);
- propagation and medium: the electromagnetic channel (`Module/UW/ElectroMagnetic/Channel`) and the MI propagation model `Module/UW/MI/CouplingPropagation` (`UwMiCouplingPropagation`).

The resulting end-to-end stack is summarised in Figure 3.1.

The two nodes are configured in a one-way link, where node 1 acts as transmitter and node 2 as receiver. The CBR application at node 1 sends fixed-size packets periodically towards node 2, which simply counts the packets received. Retransmissions are disabled (`setNoAckMode`) so that the PDR directly reflects the performance of the MI physical layer without MAC-level recovery.

3.2.2 SPECTRAL MASK AND CARRIER

DESERT's physical layer requires each signal to be associated with a spectral mask object. In the implementation, a rectangular spectral mask (`MSpectralMask/Rect`) is instantiated and attached to both `UwMiPhy` instances:

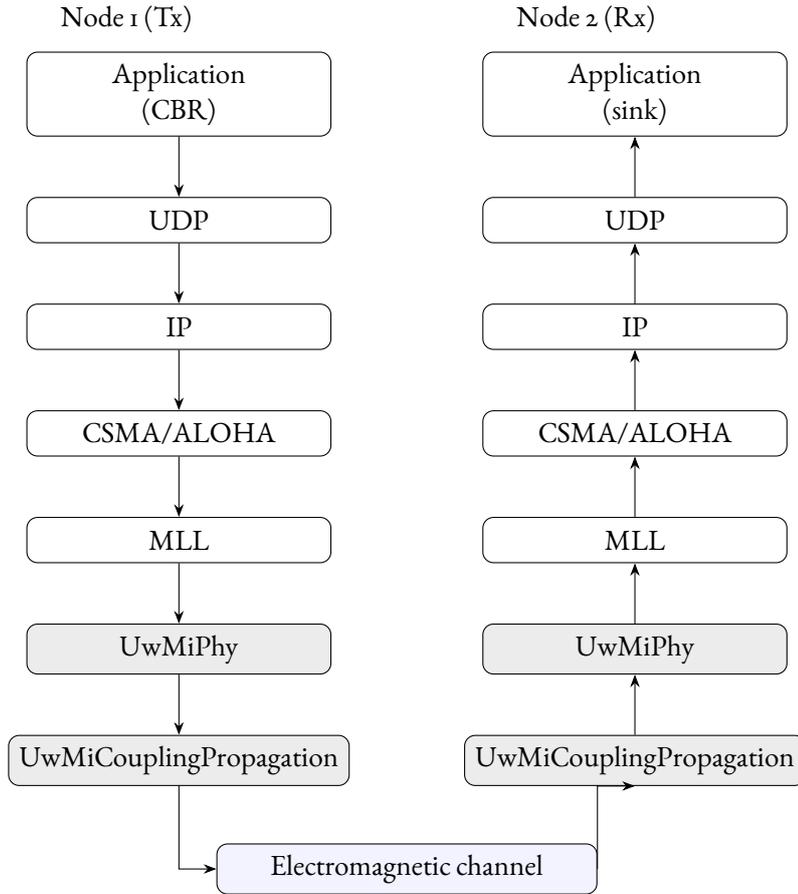


Figure 3.1: End-to-end DESERT stack used in the simulations. Node 1 (Tx) runs the standard CBR-UDP-IP-CSMA/ALOHA stack downward towards the MI PHY, while Node 2 (Rx) processes packets upward from the MI PHY to the sink application. The custom MI modules (UwMiPhy and UwMiCouplingPropagation, shaded) interface with the common electromagnetic channel.

- its centre frequency is set to $f_0 = f_{\text{khz}} \cdot 10^3$ Hz, where f_{khz} is a Tcl parameter;
- its bandwidth is set to B , another Tcl parameter.

This mask is used both for logging and, via `UwMiPhy`, in the computation of the effective noise bandwidth. Conceptually, the rectangular spectral mask plays two roles. First, it informs the DESERT channel which frequency band the MI signal occupies, so that different signals could in principle coexist in frequency. Second, its bandwidth parameter B acts as the “configurable mask bandwidth”: this is the value that `UwMiPhy` reads from the `MSpectralMask/Rect` object and uses as the starting point for the effective noise bandwidth B_{eff} . If the mask bandwidth is not set, the PHY falls back to its internal parameter

B_.

3.2.3 CONFIGURATION VIA TCL

The Tcl script `test_uwmi_net.tcl` acts as a thin but expressive configuration layer around the C++ modules. All physical, geometric and simulation parameters are exposed as command-line flags of the form `-namevalue` and stored in internal `opt(*)` variables, which are then propagated to the C++ objects through the DESERT `set` interface.

For convenience, parameters are grouped as follows:

- *Geometry and medium:* link distance and depth are controlled via `-distance` (alias `-dx`), `-z1`, `-z2` and the mode flag `-uw2aw` (`o` for `UW→UW`, `i` for `UW→AW`). The water conductivity is set by `-sigma`, with typical values $\sigma = 0.05$ S/m (fresh) and $\sigma = 4.0$ S/m (salt).
- *RF and noise parameters:* the radio-frequency (RF) carrier, bandwidth and bit rate are configured via `-f_khz`, `-B` and `-Rb`.

The quality factor Q , noise figure (NF) and noise temperature are set by `-Q`, `-NF_db` and `-Tnoise_K`, respectively, while the transmit power uses the convenient dBm interface `-txpwr_dbm`.

- *Coil and array parameters:* coil turns, radii, resistances and the coupling factor are set through `-Nt`, `-Nr`, `-at`, `-ar`, `-Rt`, `-Rr` and `-kappa`. Multi-coil stacks per node are enabled via `-Nt_coils` and `-Nr_coils`, with axial spacings controlled by `-st` and `-sr`. The flag `-auto_scale_R` switches between the “raw coil count” and “equal copper mass” regimes described in Section 3.1.
- *Simulation control:* application start/stop times, simulation stop time, packet size and CBR period are set by `-app_start`, `-app_stop`, `-sim_stop`, `-pktsize` and `-period`, while `-maxpkts` limits the maximum number of packets per run. A reproducible seed for the ns-2 pseudo-random number generator can be supplied via `-seed`.

With this arrangement it is easy to script large parameter sweeps: each point in the design space corresponds to one call to `nstest_uwmi_net.tcl` with a specific set of flags. The script also instantiates the spectral mask, connects the MI PHY and propagation modules to `Module/UW/ElectroMagnetic/Channel`, attaches position objects to the nodes, and creates the per-run logs and CSV summaries introduced in Section 3.5.

3.3 MI PROPAGATION MODEL: `UwMiCouplingPropagation`

The `UwMiCouplingPropagation` class is responsible for translating node positions, coil parameters and medium properties into a link gain that can be used by the PHY. It derives from DESERT's `MPropagation` base class and is registered under the Tcl name `Module/UW/MI/CouplingPropagation`.

3.3.1 POSITION HANDLING AND TWO-LAYER SPLITTING

During initialisation, the Tcl script creates two position objects (`Position/BM`) and attaches them to the nodes. The propagation module maintains an internal list of pointers to these positions, which is filled via the Tcl command `addPosition<PositionObj>`.

When the channel calls `getGain(Packet*)`, the module retrieves the current transmitter and receiver positions, either from its internal list or from the packet headers if necessary. If the two-layer model is enabled (`use_two_layer=1`), the helper function `splitUnderwaterAir_()` computes:

- the total distance D between the nodes;
- the length D_w of the segment lying underwater (for which conductive loss applies).

The remaining segment $D_a = D - D_w$ is implicitly treated as air.

Special corner cases, such as both nodes being entirely underwater or entirely in air, and nearly horizontal links near the interface, are handled explicitly to avoid numerical instabilities.

3.3.2 MUTUAL INDUCTANCE OF ARRAYED COILS

For a given centre-to-centre distance d , the mutual inductance between two coaxial circular coils with N_t and N_r turns, radii a_t and a_r , and medium permeability μ , is approximated in the code as

$$M(d) = \frac{\mu N_t N_r A_t A_r}{2\pi d^3}, \quad A_t = \pi a_t^2, \quad A_r = \pi a_r^2. \quad (3.6)$$

This expression corresponds to the mutual inductance of two magnetic dipoles in the near field with coaxial alignment and is implemented in `mutualInductance_onepair_()`. It follows the classical near-field dipole approximation for coaxial loops [8].

When each node hosts a stack of multiple coils, the total mutual inductance is taken as the sum of the pairwise contributions between all transmit and receive coils:

$$M_{\text{sum}}(D) = \sum_{i=1}^{N_{t,\text{coils}}} \sum_{j=1}^{N_{r,\text{coils}}} M(|D + z_j^{(r)} - z_i^{(t)}|), \quad (3.7)$$

where $z_i^{(t)}$ and $z_j^{(r)}$ are the axial offsets within each stack. The helper function `sumMutualInductance_()` implements this double sum, with a small positive lower bound on distance to avoid division by zero.

3.3.3 CONDUCTIVE LOSS AND SKIN DEPTH

To account for conductive loss in water, the propagation module uses a classical skin depth model. For a plane wave at angular frequency $\omega = 2\pi f$ in a medium with conductivity σ and permeability μ , the skin depth is

$$\delta = \sqrt{\frac{2}{\omega\mu\sigma}}. \quad (3.8)$$

as in standard conductor attenuation models [14]. In the implementation, the power attenuation over an underwater segment of length D_w is approximated as

$$L_m(f, D_w) = \exp\left(-\frac{2D_w}{\delta}\right), \quad (3.9)$$

which corresponds to an exponential decay of power with depth. This factor is computed in `conductiveFactor_(f, d)` and clipped to a minimum of 10^{-12} for numerical robustness. When `use_cond_loss_` is set to zero or $\sigma \leq 0$, the function returns 1 and the medium is effectively treated as lossless.

3.3.4 OVERALL LINK GAIN

Given the summed mutual inductance $M_{\text{sum}}(D)$ and the effective resistances $R_{t,\text{eff}}$ and $R_{r,\text{eff}}$, the propagation module computes a coupling gain

$$G_c = \kappa^2 \frac{\omega^2 M_{\text{sum}}^2}{4R_{t,\text{eff}}R_{r,\text{eff}}}, \quad (3.10)$$

where κ is a dimensionless coupling factor that captures matching and other implementation losses and gains, and $\omega = 2\pi f$. The total power gain is then

$$G = G_c L_m(f, D_w), \quad (3.11)$$

which combines the near-field geometric coupling with the conductive attenuation over the underwater segment. To avoid pathological values, G is clipped to the interval $[10^{-30}, 1]$ in the code.

The `getGain()` method returns this quantity to the physical layer, which interprets it as a linear power gain. This is the only interface through which the MI channel model affects the PHY: all further processing is performed by `UwMiPhy`.

3.4 MI PHYSICAL LAYER: `UwMiPhy`

The `UwMiPhy` class implements an analytical PER model for MI links. It derives from DESERT's `MPhy_Bpsk` and is registered under the Tcl name `Module/UW/MI/PHY/Custom`. By reusing the existing BPSK modulation machinery, it focuses on three tasks:

1. computing the received power from the transmit power and the propagation gain;
2. estimating the noise power from the noise figure and the effective bandwidth;
3. mapping the resulting SNR into E_b/N_0 , BER and PER.

3.4.1 TRANSMIT PATH

On the transmit side, `startTx()` performs the following steps:

- obtains the common and MPhy headers of the packet;
- attaches the rectangular spectral mask if none is present;
- sets the modulation type to a dedicated MI identifier, so that the receiver can recognise MI packets;
- reads the configured transmit power `TxPower_` (already expressed in Watts and set from Tcl by converting the user-level `-txpwr_dbm` argument) and stores it in `ph->Pt`;

- sets the packet duration in both the ns-2 common header and the MPhy header, using the bit rate R_b and the packet size if necessary.

For convenience, the Tcl scenario lets the user specify the transmit power in dBm through the flag `-txpwr_dbm`. The script converts this value to Watts, $P_t = 10^{(P_{\text{dBm}}-30)/10}$, and assigns the result to `TxPower_` in both `UwMiPhy` instances. As a result, the C++ code always works with linear power units and never has to perform dBm-to-Watt conversions internally.

The actual symbol-level behaviour is delegated to the base `MPhy_Bpsk::startTx()` method, ensuring compatibility with other DESERT components.

3.4.2 RECEIVE PATH AND ACQUISITION TEST

On the receive side, the DESERT channel calls `startRx()` at the beginning of a frame and `endRx()` when the frame ends. The `UwMiPhy::startRx()` method mirrors the base behaviour, accepting a packet only if the PHY is idle and no transmission is pending.

For each candidate packet, the PHY first checks that a spectral mask is attached and that the transmit power field `ph->Pt` has been initialised. It then computes:

- the received power P_r via `getRxPower()`, which multiplies the transmit power by the propagation gain returned by `UwMiCouplingPropagation`;
- the noise power N via `getNoisePower()`, which uses the effective bandwidth as described below.

The instantaneous SNR in dB is obtained as

$$\text{SNR}_{\text{dB}} = 10 \log_{10} \left(\frac{P_r}{N} \right). \quad (3.12)$$

A simple acquisition test is performed by comparing this SNR and the received power in dBm against configurable thresholds:

$$\text{SNR}_{\text{dB}} \geq \gamma_{\text{min}}, \quad P_r \geq P_{\text{thr}}, \quad (3.13)$$

where γ_{min} is set by `AcquisitionThreshold_dB_` and the power threshold P_{thr} can either be fixed or automatically derived from the noise level (if `use_auto_rx_power_gate_` is enabled). Only packets that pass this test are considered further and propagated towards the MAC.

3.4.3 EFFECTIVE BANDWIDTH AND THERMAL NOISE

A key quantity is the effective noise bandwidth B_{eff} . The PHY maintains a configurable mask bandwidth B and a resonance frequency f_0 with quality factor Q . The helper function `effectiveBandwidthHz_()` implements the following logic:

- if resonant behaviour is disabled or $Q \leq 0$, set $B_{\text{eff}} = B$;
- otherwise, compute the resonant bandwidth $B_{\text{res}} = f_0/Q$ and set $B_{\text{eff}} = \min(B, B_{\text{res}})$.

This logic is equivalent to a simple truncation rule: the effective noise bandwidth is set by the narrower of the configurable mask bandwidth B (taken from the rectangular spectral mask defined in `Tcl`) and the resonance bandwidth B_{res} . For low Q or wide coils, $B_{\text{eff}} \approx B$ and the front end behaves like an ideal rectangular filter of width B . For high Q , the resonance limits the usable band and $B_{\text{eff}} \approx B_{\text{res}} = f_0/Q$, which reflects the fact that the coil itself acts as a band-pass filter that suppresses out-of-band noise.

Given B_{eff} , the thermal noise power is evaluated as

$$N = kT_{\text{noise}}FB_{\text{eff}}, \quad (3.14)$$

where T_{noise} is the configurable equivalent noise temperature and the noise factor $F = 10^{\text{NF}_{\text{dB}}/10}$ is derived from the noise figure NF_{dB} . This computation is implemented in `thermalNoise_W_()`.

3.4.4 MAPPING SNR TO PER

The core of the PHY is the function `computePER_()`, which maps the received power P_r , the effective bandwidth and the bit rate into a PER. Given P_r and N , the linear SNR is

$$\text{SNR}_{\text{lin}} = \frac{P_r}{N}, \quad (3.15)$$

with bounds applied in the code to prevent underflow or overflow. The corresponding E_b/N_0 is then computed as

$$\left(\frac{E_b}{N_0}\right)_{\text{lin}} = \text{SNR}_{\text{lin}} \frac{B_{\text{eff}}}{R_b} 10^{-\text{NM}_{\text{dB}}/10}, \quad (3.16)$$

where NM_{dB} is an optional noise margin in dB that can be used to introduce a safety margin or to absorb unmodelled losses. The bit error probability follows from the standard BPSK formula, and the PER is obtained via the packet-length formula given earlier.

For reproducibility and post-processing, `computePER_()` emits a concise log line for each packet with tag `UWMI_METRIC` containing the main link metrics (received power, noise power, SNR, E_b/N_0 , effective bandwidth, bit rate, packet size and theoretical PER). This log is later parsed by the Python scripts that generate the PER-vs- E_b/N_0 and E_b/N_0 -vs-distance plots in Chapter 4.

The `endRx()` method finally draws a Bernoulli trial with probability P_e and marks the packet as erroneous or correct accordingly. The packet is then passed to the upper layers irrespective of the outcome; error handling (e.g. dropping corrupted packets) is left to the MAC and higher layers.

3.5 SIMULATION SCENARIOS AND LOGGING

The Tcl script `test_uwmi_net.tcl` ties together the MI PHY, the propagation model and the networking stack into a flexible simulation environment.

3.5.1 DEFAULT CONFIGURATION

Unless otherwise stated, the default configuration used in Chapter 4 is as follows:

- centre frequency $f_0 = 200$ kHz;
- bit rate $R_b = 1$ kbps;
- rectangular mask bandwidth $B = 5$ kHz;
- coil radii $a_t = a_r = 0.12$ m;
- number of turns $N_t = N_r = 40$;
- series resistances $R_t = R_r = 4 \Omega$;
- coupling factor $\kappa = 0.7$;
- transmit power configured as $P_t = 0$ dBm (i.e. 1 mW) via the `-txpwr_dbm` flag and converted to Watts in Tcl before being passed to `UwMiPhy`;
- noise figure $NF = 18$ dB;

- equivalent noise temperature $T_{\text{noise}} \approx 290$ K;
- number of coils per node $N_{t,\text{coils}} = N_{r,\text{coils}} = 1$, unless array experiments are performed.

Node positions are set according to the chosen mode (UW→UW or UW→AW) and the command–line distance parameter, and the propagation module is informed of the positions via the `addPosition` commands.

3.5.2 ONE-WAY LINK AND WATCHDOG

The CBR source at the transmitter starts at time `app_start` and sends packets of size `pktsize` bytes every `period` seconds until `app_stop` or until a maximum number of packets (if specified) has been sent. A simple Tcl watchdog procedure `stop_when_reached` periodically checks the number of packets sent and stops the application when the target is reached. The simulation itself terminates at time `sim_stop`, slightly after the application stops, to allow in–flight packets to be delivered.

3.5.3 RUN-TIME LOGGING AND CSV SUMMARIES

Two log files are maintained during each run:

- a run log (`uwmi_runlog.txt`) containing time–stamped counters of packets sent and received;
- a PHY log (`uwmi_phylog.txt`) capturing periodic debug messages from `UwMiPhy` and `UwMiCouplingPropagation`.

At the end of each simulation, a summary line is appended to a per–distance comma–separated values (CSV) file and a global master CSV file (`uwmi_results_summary_all.csv`) containing the overall packet error rate, the main configuration parameters (frequency, bandwidth, bit rate, Q , NF, transmit power, coil parameters) and the geometry (actual distance, underwater and air segment lengths, node depths, number of coils, array spacing and resistance scaling).

These CSV files, together with the per–packet `UWMI_METRIC` logs produced by the PHY, form the basis for the numerical results reported in Chapter 4. They allow traceability between each plotted point and the underlying simulation configuration and MI link budget.

3.6 EXPERIMENT AUTOMATION AND PARAMETER SWEEPS

While a single configuration of the MI link already exercises the end-to-end DESERT stack, the design questions tackled in Chapter 4 require systematic exploration of a multi-dimensional parameter space. To this end, the Tcl scenario `test_uwmi_net.tcl` is wrapped by a shell script (`sweep_runs_v2.sh`) that automates the generation of large ensembles of runs and ensures consistent logging.

The script defines a generic helper:

```
run_ns() {
  ns test_uwmi_net.tcl \
    -f_khz "$F_KHZ" -B "$B" -Rb "$RB" -Q "$Q" \
    -Nt "$NT" -Nr "$NR" -at "$AT" -ar "$AR" \
    -Rt "$RT_" -Rr "$RR_" -kappa "$KAPPA" \
    -txpwr_dbm "$TXPWR_DBM" \
    -period "$PERIOD" \
    -app_start "$APP_START" -app_stop "$APP_STOP" \
    -sim_stop "$SIM_STOP" -maxpkts "$PKTS_PER_RUN" \
    "$@"
}
```

so that any sweep can be expressed as a loop over the additional arguments "`\protect\TU\textdollar@`" (distance, conductivity, mode, array size, etc.). Timing variables `APP_START`, `APP_STOP` and `SIM_STOP` are chosen so that each run transmits a fixed number of packets (`PKTS_PER_RUN`), independent of the actual period, which simplifies the interpretation of the resulting PDR.

Several structured sweep patterns are used:

- Core distance grids: the helper `run_core_grids` sweeps the centre-to-centre distance D for the four combinations ($UW \rightarrow UW$, $UW \rightarrow AW$) \times ($\sigma = 0.05$ S/m, $\sigma = 4.0$ S/m). The distance step and maximum range are chosen according to the conductivity: fresh water grids extend farther, whereas seawater grids are shorter but use a finer spacing. Each distance point is simulated for several random seeds (set by the environment variable `SEEDS_PER_POINT`) to reduce Monte Carlo variance.
- $UW \rightarrow AW$ depth families: `sweep_uwa_depth_family` explores how performance changes with the receiver height above the surface. For a fixed σ and carrier configuration, it sweeps a list of z_2 values and, for each z_2 , runs a distance grid similar to the

core case. The result is a family of PDR-versus-distance curves and two-dimensional PDR heatmaps over (D, z_2) .

- Array gain sweeps: `sweep_arrays` controls the number of coils per node (`-Nt_coils`, `-Nr_coils`) and the axial spacing (`-st`, `-sr`). In the experiments reported later, it generates configurations with $N_{t,\text{coils}} = N_{r,\text{coils}} \in \{1, 2, 3\}$ (1×1 , 2×2 and 3×3 stacks) and two resistance regimes, selected via the flags `-auto_scale_R0` (raw coil count) and `-auto_scale_R1` (equal copper mass). Several seeds are simulated for each array size, distance and resistance regime.
- Ablation sweeps: three additional helpers, `sweep_Q`, `sweep_NF` and `sweep_B_Rb_grid`, vary one parameter at a time around the default configuration: the quality factor Q , the noise figure NF and the (B, R_b) pair. These sweeps are used to build the sensitivity plots and range tables as a function of Q , NF and effective signalling bandwidth.

A dedicated helper routine implements a range finder based on PDR. For a given conductivity σ , mode ($UW \rightarrow UW$ or $UW \rightarrow AW$), PDR target τ , frequency f_0 and bit rate R_b , the function `find_range_pdr_ge` works as follows:

1. starting from a minimum distance D_0 , it repeatedly calls `run_ns` with increasing D until the average PDR across seeds at that distance falls below τ , or a maximum distance is reached (bracketing phase);
2. it then performs a binary search between the last “good” distance and the first “bad” one, again using the aggregated PDR from the master CSV, and returns the largest D for which $\text{PDR} \geq \tau$.

This routine is used, for example, to build the range-versus-frequency and range-versus-parameter curves in Chapter 4 using $\tau = 0.8$ or $\tau = 0.9$ as typical reliability targets. Each point on those curves therefore corresponds to a cluster of DESERT runs whose configuration is fully traceable through the CSV and log filenames.

3.7 POST-PROCESSING, CONFIDENCE INTERVALS AND RANGE EXTRACTION

The raw outputs of the DESERT runs are processed offline by a Python script (`make_core_plots.py`) that consolidates the per-run CSV summaries and the per-packet PHY logs into a coherent set of curves and design tables. The processing pipeline is deliberately simple and reproducible, and mirrors the system model introduced earlier in this chapter.

3.7.1 AGGREGATION OF PDR AND WILSON INTERVALS

Each simulation run appends one line to the master CSV `uwmi_results_summary_all.csv`, including the number of packets sent and received, the packet error rate, the main physical parameters (f_0 , B , R_b , Q , NF, transmit power, coil geometry) and the geometry (actual distance D , underwater and air segment lengths D_w , D_a , node depths z_1 , z_2 , array parameters and resistance regime).

The Python script first groups runs by a key

$$(D, \sigma, \text{mode}, f_0, R_b, B),$$

where “mode” is $UW \rightarrow UW$ or $UW \rightarrow AW$. For each such group it accumulates the total number of packets sent and received across all seeds and computes an empirical packet delivery ratio (PDR),

$$\widehat{\text{PDR}} = \frac{\sum_i \text{rcv}_i}{\sum_i \text{sent}_i}.$$

To quantify the uncertainty due to finite sampling, it then computes a binomial Wilson score interval at 95 % confidence for each point, yielding lower and upper bounds (PDR_{lo} , PDR_{hi}). These aggregated quantities are used consistently in Chapter 4: all PDR-versus-distance plots display $\widehat{\text{PDR}}$ with a shaded band indicating the corresponding Wilson interval. In practice, each aggregated point pools packets over multiple seeds, so the confidence intervals are dominated by channel behaviour rather than Monte Carlo noise.

The same aggregated table underlies the range estimates. Given a target reliability τ (typically $\tau = 0.8$ or 0.9) and a fixed configuration $(\sigma, \text{mode}, f_0, R_b, B)$, the script extracts the PDR sequence as a function of distance and computes the largest distance D at which $\widehat{\text{PDR}} \geq \tau$. If the PDR crosses τ between two sampled distances, a linear interpolation in $(D, \widehat{\text{PDR}})$ is applied. This definition is implemented in a helper function `range_from_P()` and is reused for all range-versus-parameter tables (e.g. range versus frequency, versus Q , versus NF, and in rate-range frontiers).

3.7.2 EXTRACTION OF PHY METRICS FROM LOGS

In addition to the master CSV, each run produces a log file in which the PHY emits compact lines prefixed by the tag `UWMI_METRIC`. Each such line contains, for a single packet, the received power in dBm, the noise power in dBm, the instantaneous SNR and the resulting E_b/N_0 in dB, together with the effective bandwidth and the packet length used in the PER

computation.

The post-processing script scans all log files in the sweep directory and parses these lines using a regular expression. Each packet is associated with its configuration and distance by decoding the log filename and, when needed, by joining on the distance and parameter fields present in the master CSV. For each combination of $(D, \sigma, \text{mode}, f_0, R_b)$ the script then computes:

- the median and interquartile range of E_b/N_0 in dB;
- the median and interquartile range of the received power in dBm.

These statistics are used to produce median E_b/N_0 -versus-distance curves with 25th/75th percentile bands, and analogous received-power-versus-distance plots. On the latter, an additional reference slope proportional to $1/d^6$ is overlaid and anchored at the shortest distance point, illustrating the $1/d^6$ decay predicted by the ideal near-field MI dipole model.

For each plotted slice, the script also reconstructs the effective noise floor used by the PER model. It does so by taking the median values of T_{noise} , NF, B and (f_0, Q) from the CSV entries belonging to that slice, rebuilding the effective bandwidth B_{eff} according to Section 3.4, and evaluating

$$N = kT_{\text{noise}}FB_{\text{eff}}$$

with $F = 10^{\text{NF}_{\text{dB}}/10}$. The resulting noise power in dBm is shown as a horizontal dashed line on the received-power plots, aligning the visualisation with the analytical noise model implemented in UwMiPhy.

3.7.3 THEORY-VERSUS-SIMULATION PER CURVES AND DESIGN TABLES

Finally, the script compares the analytical BPSK PER model with the simulated MI link. Using the dominant packet size observed in the CSV, it first builds a theoretical PER curve $P_e(E_b/N_0)$ for uncoded BPSK in AWGN. It then pairs, for each configuration and distance, the empirical PER from the CSV with the median E_b/N_0 recovered from the UWMI_METRIC logs. These pairs are aggregated across distances into a set of points plotted on the PER-versus- E_b/N_0 plane, demonstrating that the simulation follows the analytical model within the expected Monte Carlo variability.

The same infrastructure is reused to generate design-oriented summaries. By applying the `range_from_P()` operator on different slices of the aggregated PDR table, the script produces:

- range-versus-frequency curves for sea water and fresh water, in both UW→UW and UW→AW modes;
- range-versus-quality-factor and range-versus-noise figure tables, quantifying the benefit of improving Q or NF for the chosen MI geometry;
- rate-range frontier tables over a grid of (B, R_b) pairs, indicating, for each signalling rate, the maximum distance at which the target PDR is still achievable;
- range-versus-array-size summaries that condense the impact of 1×1 , 2×2 and 3×3 coil stacks under both “raw copper” and “equal copper mass” regimes.

Together with the C++ and Tcl code described earlier, these scripts form a simple but complete pipeline that links the physical MI model to packet-level metrics and design tables, while keeping each plotted point traceable to the underlying DESERT runs.

3.8 CHAPTER SUMMARY

This chapter has detailed the system model and the DESERT implementation used to study underwater MI communication in the remainder of the thesis. The core contributions are:

- a propagation model that combines a near-field MI coupling law for coil arrays with a two-layer conductive loss model for UW→UW and UW→AW links;
- a physical layer that converts the MI link budget into analytical BPSK error probabilities in AWGN, while remaining fully compatible with DESERT’s modular architecture;
- a Tcl-driven simulation environment that exposes the main hardware and geometry parameters and logs all relevant metrics in machine-readable form, together with automated sweep and post-processing scripts that turn these logs into statistically robust performance curves and design tables.

Together, `UwMiCouplingPropagation`, `UwMiPhy` and the associated scripts provide a physically grounded yet computationally efficient platform for exploring how conductivity, frequency, coil geometry and hardware design choices shape the achievable range and reliability of underwater MI links. Because all key parameters are exposed via the Tcl interface and all results are logged in machine-readable CSV form, this platform can be reused as a small simulation *framework* for other DESERT users who wish to study different MI modem designs without modifying the core propagation or PHY code. The next chapter leverages this platform to produce the numerical results and design guidelines presented in Chapter 4.

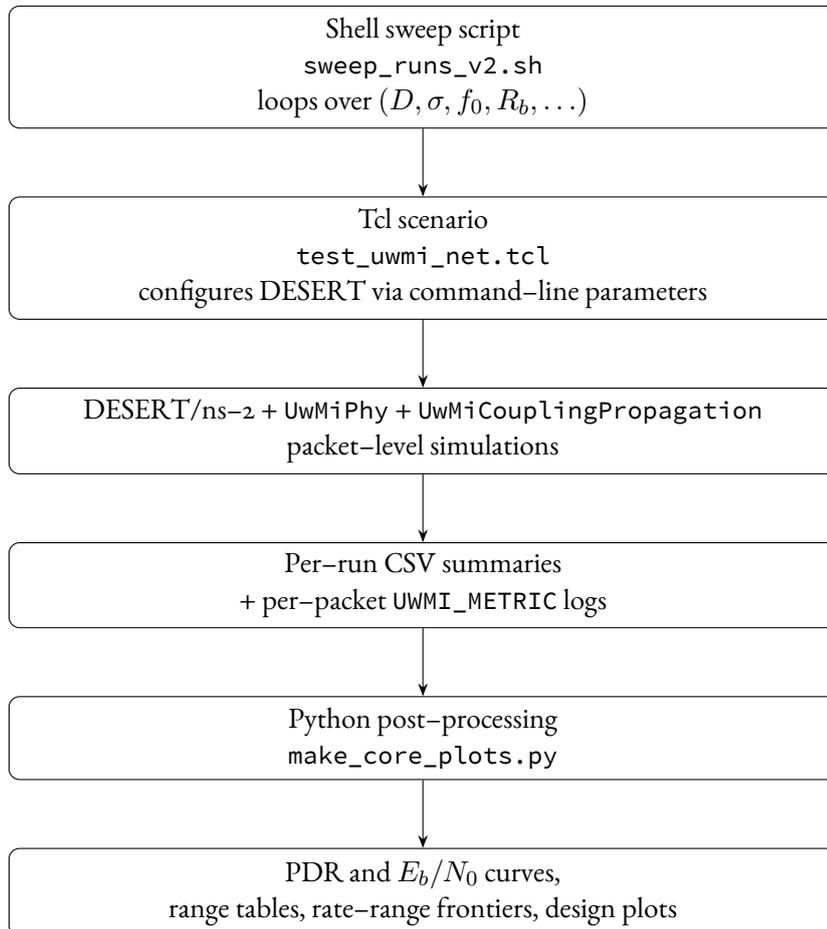


Figure 3.2: Automation pipeline for MI simulations. The shell script `sweep_runs_v2.sh` sweeps the design space and repeatedly calls the Tcl scenario `test_uwmi_net.tcl`, which configures DESERT/ns-2 runs with `UwMiPhy` and `UwMiCouplingPropagation`. Each run produces CSV summaries and per-packet `UWMI_METRIC` logs, which are parsed by the Python script `make_core_plots.py` to generate PDR curves, E_b/N_0 metrics and range tables used in Chapter 4.

“In underwater magneto–inductive systems, physics dictates the outcome: magnetic near–field coupling decays steeply with distance, and water conductivity sets the limit.”

4

Results

This chapter presents the numerical evaluation of the magneto–inductive underwater communication model introduced in Chapter 3. All simulations are carried out in the DESERT Underwater framework [10], using the custom `UwMiCouplingPropagation` and `UwMiPhy` modules developed in this work.

The main goals of the evaluation are:

- to validate the proposed PHY model against analytical BPSK performance in AWGN;
- to characterise the achievable range of a single–coil MI link in representative freshwater and seawater conditions, both for underwater–to–underwater (UW→UW) and underwater–to–air (UW→AW) geometries;
- to quantify the impact of node depth and of the water–air interface on link reliability;
- to assess the gain provided by multi–coil transmit and receive arrays, under both “more copper” and “equal copper mass” design constraints;
- to explore the sensitivity of the link range to key hardware parameters (centre frequency f_0 , quality factor Q , receiver noise figure and system bandwidth), and to derive rate–range trade–offs.

Unless otherwise stated, results are obtained at $f_0 = 200$ kHz, bit rate $R_b = 1$ kbps, rectangular spectral mask of $B = 5$ kHz, and BPSK modulation over an equivalent AWGN channel. The default configuration uses single circular coils ($N_t = N_r = 40$ turns, radius

$a_t = a_r = 0.12$ m, resistances $R_t = R_r = 4 \Omega$, coupling factor $\kappa = 0.7$) and transmit power $P_t = 0$ dBm. Two conductivities are considered: $\sigma = 0.05$ S/m (clean or fresh water) and $\sigma = 4.0$ S/m (salt water).

These baseline parameters are chosen to approximate a simple, low-cost MI modem: a 12 cm radius coil with 40 turns and $\approx 4 \Omega$ series resistance, driven at 0 dBm transmit power and a moderate receiver NF around 18 dB. They are of the same order of magnitude as small MI prototypes reported in the literature [11, 12, 13], so the ranges reported below can be interpreted directly for low-power underwater sensors and robots using inexpensive hardware.

SIMULATION FRAMEWORK AND AUTOMATION

All results in this chapter are obtained with the MI simulation framework introduced in Chapter 3. In particular, the Tcl scenario `test_uwmi_net.tcl` configures the DESERT stack with `UwMiPhy` and `UwMiCouplingPropagation` (Section 3.2), while the shell script `sweep_runs_v2.sh` wraps this scenario to sweep the design space over distance, conductivity, carrier frequency, number of coils, quality factor, noise figure and rate (Section 3.6). Each parameter point is simulated for multiple random seeds, and the resulting per-run CSV summaries and per-packet `UWMI_METRIC` logs are post-processed by the Python script `make_core_plots.py` as described in Section 3.7.

In practice, every marker in the figures below corresponds to an aggregate of many packet-level runs. The post-processing stage pools statistics across seeds, computes Wilson confidence intervals for the PDR, and extracts range metrics D_τ by interpolation on the PDR-versus-distance curves. Because all physical and system parameters are exposed as Tcl options and all outputs are stored in machine-readable CSV form, the plots and tables in this chapter can be reproduced by re-running the same sweeps with modified configuration flags, without changing the underlying C++ or Tcl code.

4.1 PHY VALIDATION IN AWGN

Before studying range, the PHY model must be checked against analytical BPSK performance. Figure 4.1 compares the packet error rate (PER) predicted by the analytical AWGN model with the empirical PER obtained from DESERT simulations. The theoretical curve is computed for uncoded BPSK over AWGN with packet length $L = 128$ bytes, using the standard expression from digital communications theory [15]. The simulated points are obtained by forcing a flat MI channel (no conductive loss, unity gain) and by sweeping the

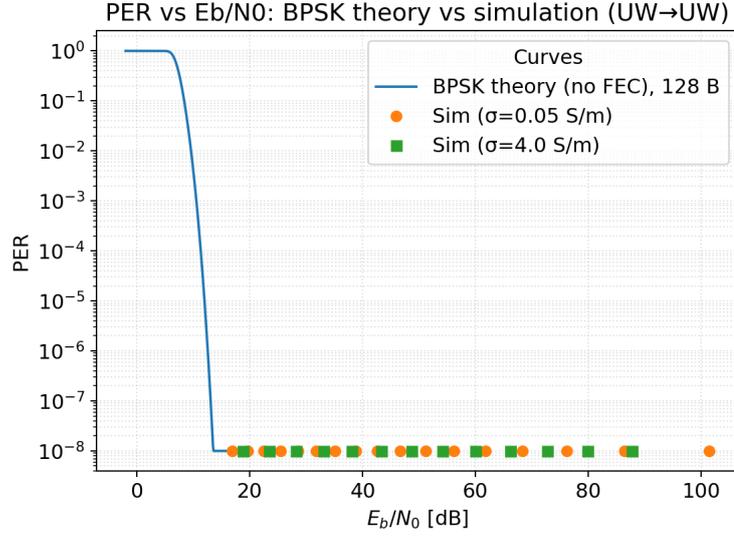


Figure 4.1: PER as a function of E_b/N_0 for uncoded BPSK: analytical curve and DESERT simulations using UwM*i*Phy. Each marker aggregates multiple seeds at the same E_b/N_0 ; PER is plotted on a logarithmic scale.

transmit power so as to cover a wide range of E_b/N_0 .

As in Chapters 2 and 3, the PHY assumes uncoded BPSK over an equivalent AWGN channel. The role of the MI propagation model is to map coil geometry and conductive loss into an equivalent E_b/N_0 , after which UwM*i*Phy behaves as a standard BPSK modem in AWGN.

The simulation points for both conductivities essentially lie on top of the analytical curve in the waterfall region. The transition from almost sure packet loss ($\text{PER} \approx 1$) to error-free operation ($\text{PER} < 10^{-8}$) occurs around $E_b/N_0 \approx 16\text{--}18$ dB, in line with the theoretical prediction for BPSK with packets of this length. At higher E_b/N_0 the PER saturates at the floor imposed by the finite number of simulated packets. Importantly, the PER depends only on E_b/N_0 and not on the medium conductivity: once the MI link budget has been converted into E_b/N_0 by UwM*i*Phy, the resulting error statistics match the AWGN reference. This validates the thermal-noise and PER computation in the custom PHY.

4.2 SINGLE-COIL LINK PERFORMANCE

4.2.1 UW→UW AND UW→AW LINKS

Figure 4.2 and Figure 4.3 report the packet delivery ratio ($\text{PDR} = 1 - \text{PER}$) as a function of the Tx–Rx distance D for a single-coil link. The default geometry for underwater-to–

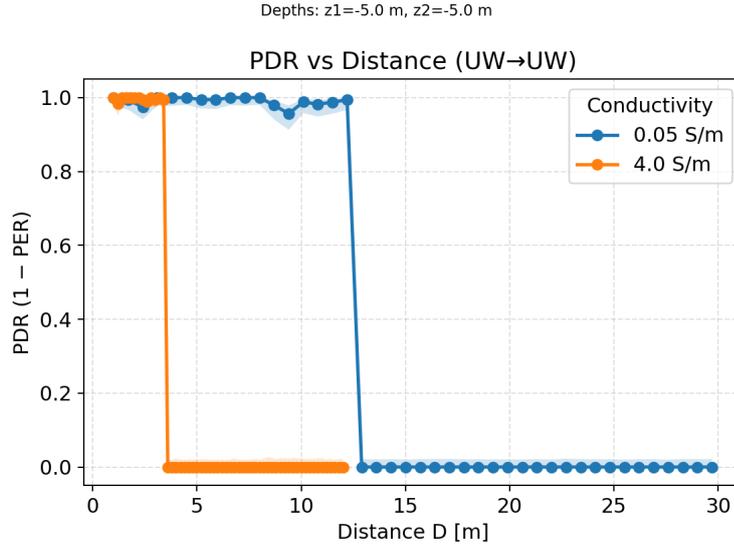


Figure 4.2: PDR vs distance for the UW→UW link in freshwater ($\sigma = 0.05$ S/m) and seawater ($\sigma = 4.0$ S/m). Depths: $z_1 = z_2 = -5$ m.

underwater operation places both nodes at $z_1 = z_2 = -5$ m, whereas the underwater-to-air configuration uses $z_1 = -2$ m and $z_2 = +1$ m. For each point, multiple independent random seeds are simulated (typically 20 realisations per distance), and the PDR is aggregated using Wilson confidence intervals.

In all cases the PDR exhibits a very sharp “waterfall”: for short distances the link is essentially error-free ($\text{PDR} \approx 1$), while beyond a critical distance the PDR collapses to zero within one or two grid steps. For freshwater ($\sigma = 0.05$ S/m) the critical distance is around 12 m for the UW→UW configuration and about 15 m for UW→AW. In seawater ($\sigma = 4.0$ S/m) the range shrinks to approximately 3.5 m and 4.5 m respectively. These ranges are summarised more precisely in Table 4.1 in terms of the maximum distance such that $\text{PDR} \geq \tau$.

The corresponding E_b/N_0 and received power are shown in Figures 4.4–4.7. For each distance the logs produced by UWMiPhy are used to compute the median E_b/N_0 and P_r across seeds.

In freshwater the link operates with a comfortable SNR margin up to the edge of coverage: at the largest distances where the PDR is still high, the median E_b/N_0 remains well above the 10 dB design target (around 17–20 dB), consistent with the nearly error-free PDR in Figures 4.2–4.3. In seawater the same transmitted power yields significantly lower E_b/N_0 due to conductive loss, and the 10 dB threshold is reached at much shorter distances.

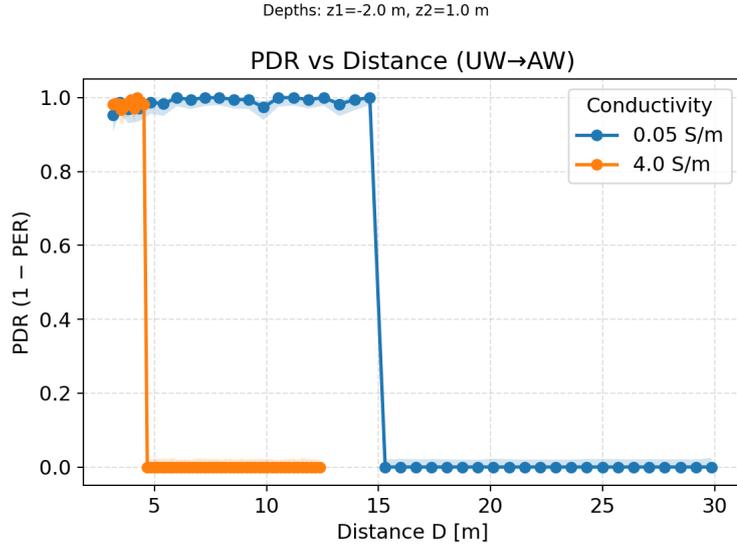


Figure 4.3: PDR vs distance for the UW→AW link in freshwater and seawater. Depths: $z_1 = -2.0$ m, $z_2 = +1.0$ m.

The Rx-power plots confirm that the implemented propagation model follows the expected magnetic-dipole law. The dotted reference curve corresponds to an ideal $P_r \propto 1/d^6$ behaviour, which is the textbook near-field scaling of a coaxial MI link. The simulated curves track this slope closely, especially at shorter distances where conductive loss is less dominant. At larger distances the curves bend downward more steeply than $1/d^6$ in seawater, due to the exponential attenuation factor derived from the skin depth. In all cases the waterfall in PDR occurs when the received signal approaches the thermal noise floor marked in Figures 4.6 and 4.7. This coherent picture across PDR, E_b/N_0 and P_r validates the link-budget logic inside `UwMiPhy` and `UwMiCouplingPropagation`.

4.2.2 RANGE AT PDR THRESHOLDS

To provide compact design numbers, the PDR curves are converted into a range metric defined as the maximum distance D_τ such that $\text{PDR} \geq \tau$. Table 4.1 reports D_τ for $\tau \in \{0.8, 0.9\}$, for both geometries and conductivities, at $f_0 = 200$ kHz and $R_b = 1$ kbps.

Because the PDR transitions are extremely steep, the range is almost independent of whether τ is set to 0.8 or 0.9. The main effect of the water-air interface is to slightly extend the range compared to a fully submerged link with the same total distance: the UW→AW configuration achieves about 2 m extra range in both media. This is consistent with the two-layer propagation model implemented in `UwMiCouplingPropagation`, where only the under-

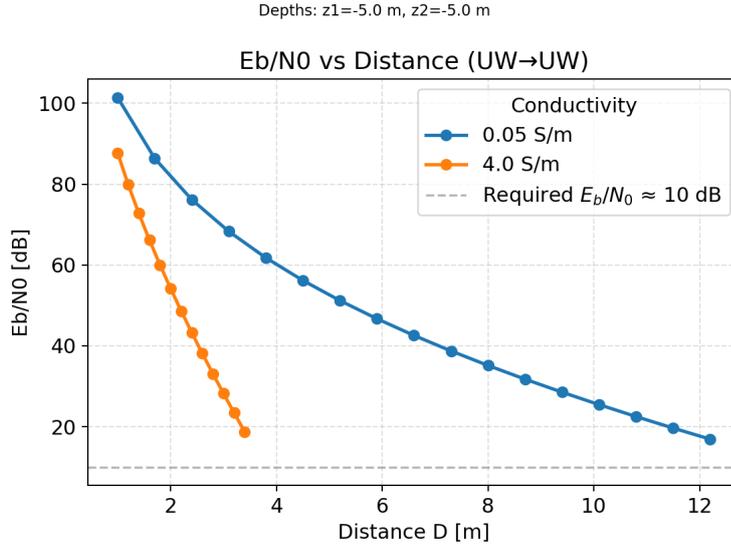


Figure 4.4: Median E_b/N_0 vs distance for the UW→UW configuration at $\sigma = 0.05$ S/m and $\sigma = 4.0$ S/m. The dashed line indicates a conservative target $E_b/N_0 \approx 10$ dB for reliable BPSK reception.

Table 4.1: Range D_τ (in metres) for different PDR thresholds τ , link geometries and conductivities. Values are obtained by interpolation on the PDR vs distance curves.

τ	Mode	σ [S/m]	D_τ [m]
0.8 / 0.9	UW→UW	0.05	12.2
	UW→UW	4.0	3.4
	UW→AW	0.05	14.6
	UW→AW	4.0	4.5

water segment undergoes conductive loss while the magnetic field in air is essentially lossless over these distances.

4.3 IMPACT OF DEPTH AND WATER–AIR INTERFACE

The previous section fixed the node depths. In practice, especially for UW→AW links, the height of the above–water node relative to the surface can vary. To investigate this effect, the UW→AW scenario is repeated while sweeping the receiver depth z_2 between 0.2 and 2.0 m above the surface, at fixed transmitter depth $z_1 = -2$ m.

Figure 4.8 shows the resulting PDR vs distance for seawater ($\sigma = 4.0$ S/m), while Figure 4.9 reports the same curves for freshwater. Each coloured curve corresponds to a different

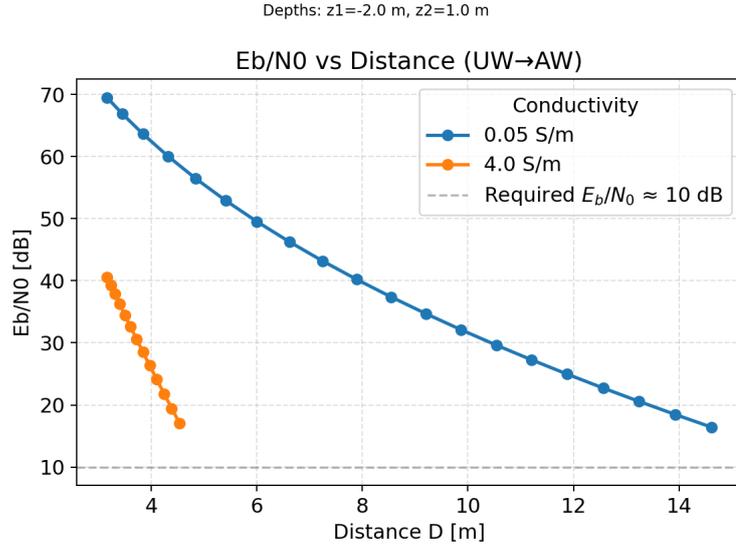


Figure 4.5: Median E_b/N_0 vs distance for the UW→AW configuration.

z_2 .

In both media, varying the receiver height over the range 0.2–2 m produces only a modest change in the coverage distance. In seawater the critical distance varies by less than one metre across all z_2 , while in freshwater the spread is on the order of one to two metres. Intuitively, once the segment in air becomes comparable to the underwater segment, the additional gain of moving the receive coil slightly higher is small.

The combined effect of distance and z_2 can be visualised as heatmaps of the PDR, shown in Figures 4.10–4.11. Yellow regions correspond to almost perfect reliability, while dark regions indicate PDR close to zero.

In freshwater the good-PDR region forms an almost vertical band up to $D \approx 15$ –16 m for all z_2 , then collapses rapidly. In seawater the band is compressed to $D \approx 4$ –5 m, again with little dependence on z_2 . From a design perspective, this means that the UW→AW MI link is quite robust to the exact placement of the above-water node: once the receive coil is a short distance above the interface, the underwater segment—and thus the conductive loss—is essentially fixed, and further increasing z_2 only changes the total distance in the $1/d^6$ near-field term. Fine tuning the buoy height within 0.2–2 m yields only modest (order-of-metres) changes in range compared to controlling distance, conductivity and coil parameters.

Overall, distance and conductivity dominate the MI link budget in the considered scenarios; depth variations of a couple of metres around the surface only introduce second-order

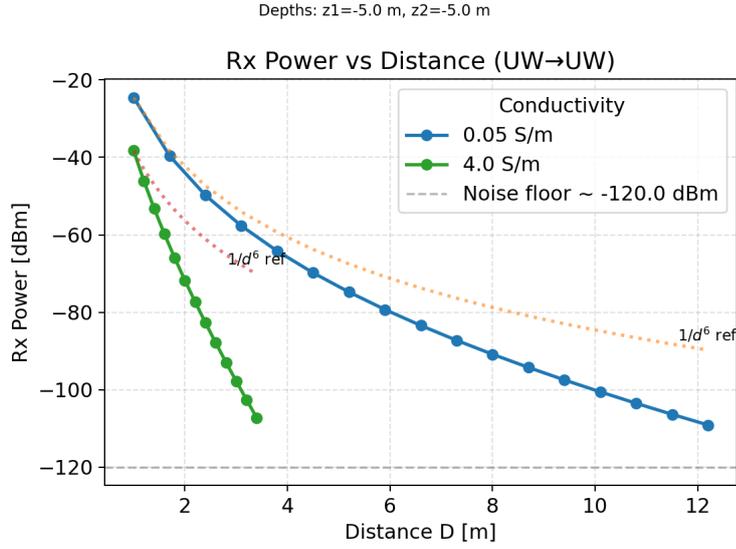


Figure 4.6: Median received power as a function of distance for the UW→UW configuration. The dotted line shows an ideal $1/d^6$ near-field MI decay, anchored at the first point. The horizontal dashed line indicates the thermal-noise floor computed from the noise figure and effective bandwidth.

corrections.

4.4 ARRAY GAIN WITH MULTI-COIL ANTENNAS

We now investigate the benefits of using multiple coils per node. In all cases the link is UW→UW in freshwater ($\sigma = 0.05$ S/m), with both nodes at $z_1 = z_2 = -5$ m and centre frequency $f_0 = 200$ kHz.

In the implementation of `UwMiCouplingPropagation`, an “array” is modelled as N identical circular coils per node, connected in series and placed along the line joining the node centres. The coil centres are offset by a fixed spacing of 0.04 m, so that for a given node the axial positions are $\{-(N-1)s/2, \dots, 0, \dots, (N-1)s/2\}$ with $s = 0.04$ m. The total mutual inductance used in the link gain is the sum of the pairwise contributions over all transmitter and receiver coils, i.e., $\sum_{i,j} M(d_{ij})$, where d_{ij} is the distance between the i -th transmit and j -th receive coil along the axis. This is therefore a one-dimensional coaxial array model; lateral displacements and tilt are not considered.

Two design constraints are explored by means of the `auto_scale_R` parameter:

1. *Raw coil count (more copper)*: the effective series resistance of each node remains equal to that of a single coil ($R_{t,\text{eff}} = R_t$, $R_{r,\text{eff}} = R_r$), so adding coils increases the amount

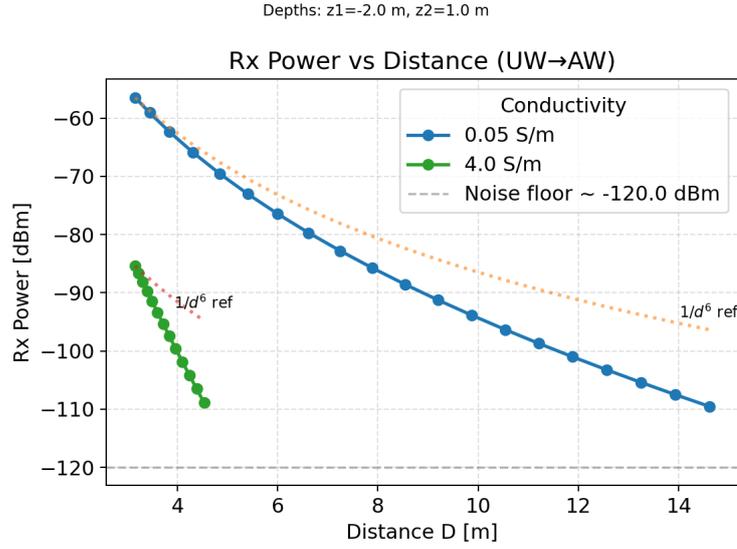


Figure 4.7: Median received power vs distance for the UW→AW configuration.

of copper and hence the mutual coupling without increasing ohmic loss;

2. *Equal copper mass*: the effective series resistance is scaled approximately linearly with the number of coils, $R_{t,\text{eff}} \approx NR_t$ and $R_{r,\text{eff}} \approx NR_r$, mimicking a design in which the total copper mass is kept roughly constant while it is split across more, thinner coils.

In both cases we consider $N \in \{1, 2, 3\}$ coils per node. These two regimes correspond exactly to the resistance-scaling options introduced in Section 3.1, and they illustrate how `auto_scale_R` can be used as a convenient high-level switch between “more copper” and “equal copper mass” array designs within the same simulation framework.

4.4.1 PDR VS DISTANCE

Figure 4.12 reports the PDR vs distance for $N = 1, 2, 3$ in the “raw coil count” setting, whereas Figure 4.13 shows the same curves under the equal-mass constraint.

In both configurations the main effect of increasing N is a horizontal shift of the waterfall region towards larger distances. With unconstrained copper, moving from a single coil to three coils per node pushes the PDR drop from about 12 m to roughly 17 m. Under the equal-mass constraint the gain is slightly smaller, with the $N = 3$ case reaching about 15 m.

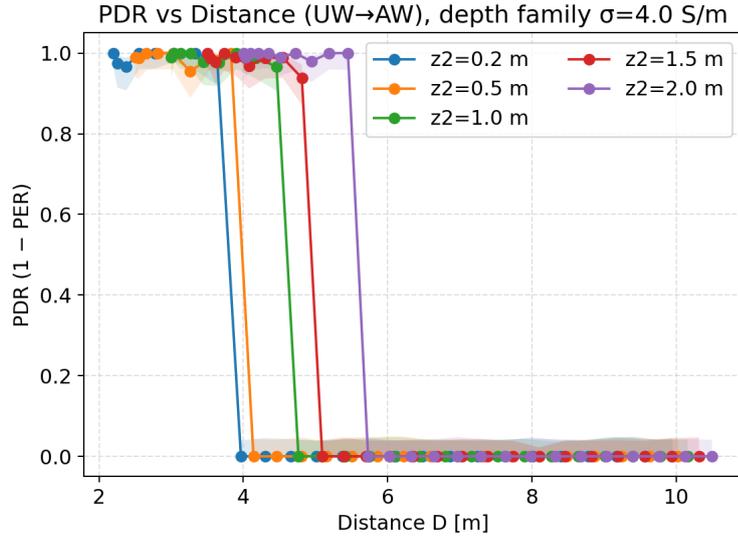


Figure 4.8: PDR vs distance for UW→AW in seawater, for different receiver heights above the surface z_2 . Conductivity $\sigma = 4.0$ S/m, $z_1 = -2$ m.

The curves remain very steep, confirming that the arrays primarily extend the usable range rather than gradually degrading reliability.

4.4.2 RANGE VS NUMBER OF COILS

The PDR curves can again be condensed into a range metric at $\text{PDR} \geq 0.8$. Figure 4.14 and Figure 4.15 summarise the resulting range as a function of N for the two design constraints.

For a single coil the range is about 12 m, in line with Table 4.1. With more copper the $N = 2$ and $N = 3$ stacks extend the range to roughly 14 m and 17 m, respectively. When the total series resistance is increased proportionally to the number of coils, the gain is reduced but still substantial, reaching about 14–15 m for $N = 3$. This behaviour is consistent with the arrays increasing the effective magnetic moment of the antenna, hence the mutual inductance, while the equal-mass constraint partly offsets this by raising the ohmic loss. It should also be emphasised that this is a coaxial one-dimensional array model; the impact of three-dimensional layouts and random misalignment is left for future work.

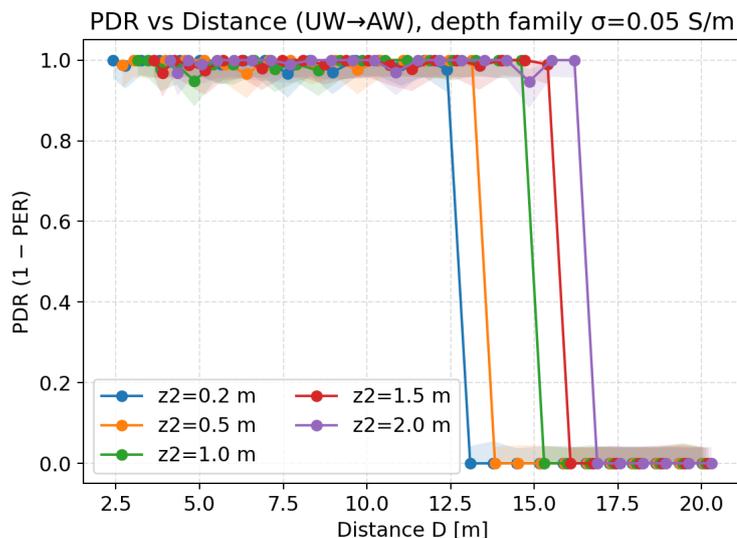


Figure 4.9: PDR vs distance for UW→AW in freshwater, for different z_2 .

4.5 SENSITIVITY TO FREQUENCY AND HARDWARE PARAMETERS

We finally study how the link range changes with the carrier frequency f_0 , the coil quality factor Q , the receiver noise figure (NF) and the effective system bandwidth.

4.5.1 RANGE VS CARRIER FREQUENCY

The first experiment sweeps f_0 between 20 and 500 kHz while keeping all other parameters fixed. Range is measured at $\text{PDR} \geq 0.8$ for both UW→UW and UW→AW links. The resulting curves are shown in Figures 4.16–4.17.

In freshwater the range is remarkably flat over almost two decades of frequency: it stays between roughly 13 and 15 m for both geometries. This reflects the fact that, in low-conductivity water, the near-field MI coupling and the coil parameters dominate over conductive loss, so changing f_0 mainly shifts the resonance without radically altering the link budget.

In seawater the picture is different. Here the skin depth becomes comparable with the link distance, and the conductive loss introduces an exponential attenuation factor that grows with frequency. As a result the range decreases monotonically with f_0 , from about 6–7 m at 20 kHz down to around 3–4 m at 300–500 kHz. This confirms the intuition that, in highly conductive media, MI communications should operate at the lowest practical frequencies.

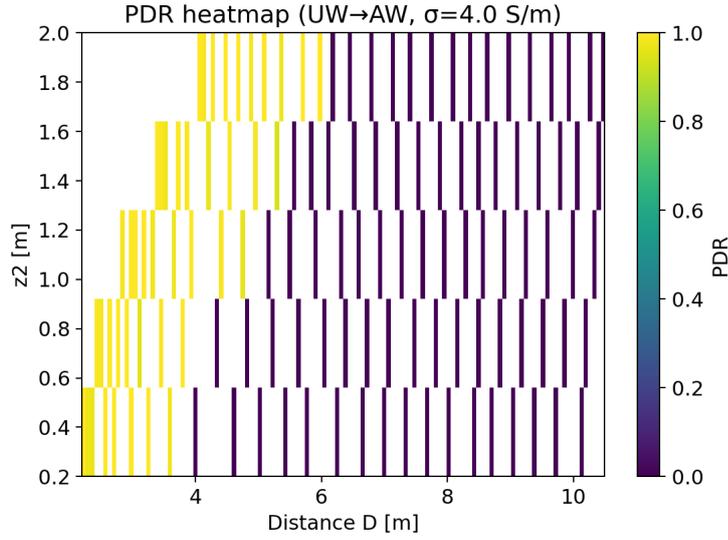


Figure 4.10: Heatmap of PDR as a function of distance D and receiver height z_2 for the UW→AW link in seawater.

4.5.2 RANGE VS COIL QUALITY FACTOR

Next, the coil quality factor Q is varied while all other parameters, including the effective noise bandwidth, are kept fixed. Range is again measured for the UW→UW freshwater link at $\text{PDR} \geq 0.8$. Table 4.2 summarises the results (values are rounded to 0.1 m).

Within the explored range of Q the achievable distance varies between about 12.1 and 13 m, with a very shallow maximum around $Q \approx 50$. This indicates that, once the centre frequency and bandwidth have been fixed, moderate variations in Q mainly perturb the matching and the effective noise bandwidth but do not change the range dramatically. In other words, the link is relatively tolerant to coil losses, which is encouraging for practical implementations.

4.5.3 RANGE VS RECEIVER NOISE FIGURE

A similar sweep is performed over the receiver noise figure, with NF ranging between 10 and 26 dB. The corresponding ranges for the UW→UW freshwater link are reported in Table 4.3.

As expected, the dependence on NF is weak over this range: all values cluster around 12–13 m, with differences on the order of 1 m, comparable with the resolution of the range estimator based on discrete distance samples. Improving the noise figure alone therefore yields only limited range gains under the present configuration.

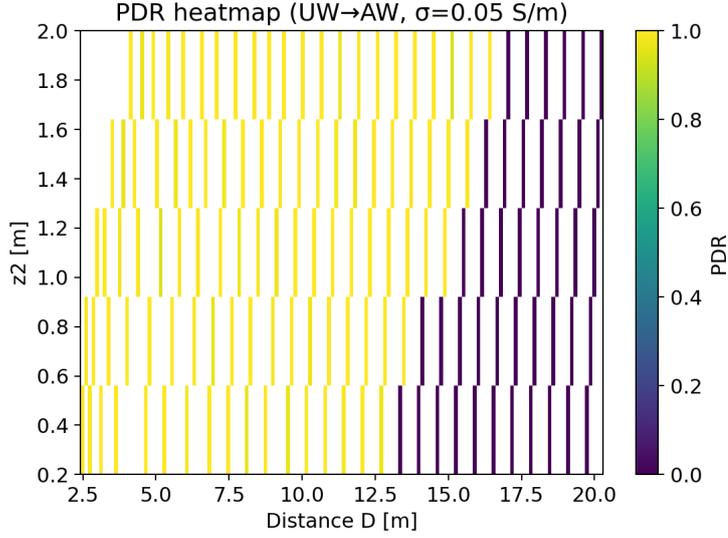


Figure 4.11: Heatmap of PDR as a function of distance D and receiver height z_2 for the UW→AW link in freshwater.

4.5.4 RATE–RANGE FRONTIER OVER BANDWIDTH AND BIT RATE

The previous sweeps varied one parameter at a time. As a compact design summary, this subsection explores the joint effect of occupied bandwidth B and bit rate R_b on the achievable range. We focus on the UW→UW link in freshwater ($\sigma = 0.05$ S/m) at $f_0 = 200$ kHz and $Q = 50$. For each pair (B, R_b) , the range metric $D_{0.8}$ is defined as the maximum distance such that $\text{PDR} \geq 0.8$. Table 4.4 reports $D_{0.8}$ for all (B, R_b) pairs considered, derived from the CSV output of the simulation sweeps.

Two observations follow. First, for all tested configurations the range stays between 12.2 and 13.2 m: doubling the bit rate from 1 to 2 kbps has practically no impact on coverage, and increasing B from 2 to 8 kHz shortens $D_{0.8}$ by about one metre. This is consistent with the effective–bandwidth expression introduced in Chapter 2: with $Q = 50$ at $f_0 = 200$ kHz the resonant bandwidth is $f_0/Q \approx 4$ kHz, so for $B \gtrsim 4$ kHz the noise bandwidth is limited by the coil rather than by the rectangular spectral mask.

Second, the small spread of $D_{0.8}$ across (B, R_b) confirms that, under the chosen hardware parameters, the MI link is primarily range–limited by medium conductivity and coil geometry. Moderate changes in spectral efficiency only provide second–order adjustments. This rate–range frontier complements the previous sweeps by showing a broad operating region where one can trade bandwidth for bit rate with negligible loss of coverage.

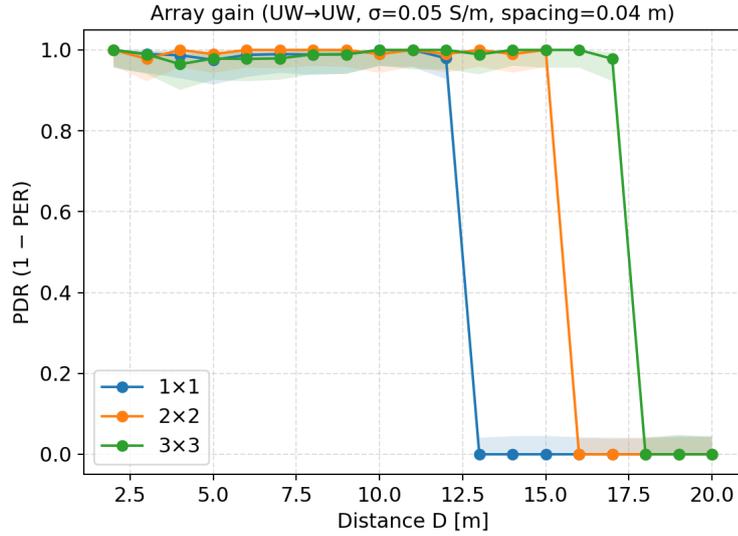


Figure 4.12: Array gain for UW→UW in freshwater ($\sigma = 0.05$ S/m): PDR vs distance for different numbers of coils per node in the “raw coil count” configuration (`auto_scale_R=0`).

4.6 DISCUSSION AND CHAPTER SUMMARY

This chapter has characterised the performance of the proposed magneto–inductive PHY layer and propagation model under a variety of underwater scenarios.

First, the AWGN validation in Section 4.1 showed that the simulated PER vs E_b/N_0 follows the analytical BPSK curve almost exactly, for both conductivities. Once the MI link budget is converted into E_b/N_0 inside `UwMiPhy`, the error statistics match the textbook AWGN reference. This gives confidence that the noise modelling and PER computation in the custom PHY are correct.

Second, the single–coil experiments quantified the baseline range of the link for representative freshwater and seawater conditions. At $R_b = 1$ kbps and $P_t = 0$ dBm, a single circular coil reaches about 12 m in UW→UW operation and about 14–15 m in UW→AW, while in seawater the range shrinks to 3–4.5 m. The PDR waterfalls, E_b/N_0 and received–power plots are mutually consistent: the waterfall occurs when the received signal approaches the thermal–noise floor and the E_b/N_0 drops below the BPSK threshold. The received power follows the expected $1/d^6$ near–field decay, with additional attenuation in seawater due to finite skin depth. This validates both the coupling model and the two–layer conductive–loss implementation of `UwMiCouplingPropagation`.

Third, the depth and interface study in Section 4.3 showed that, for UW→AW links, vary-

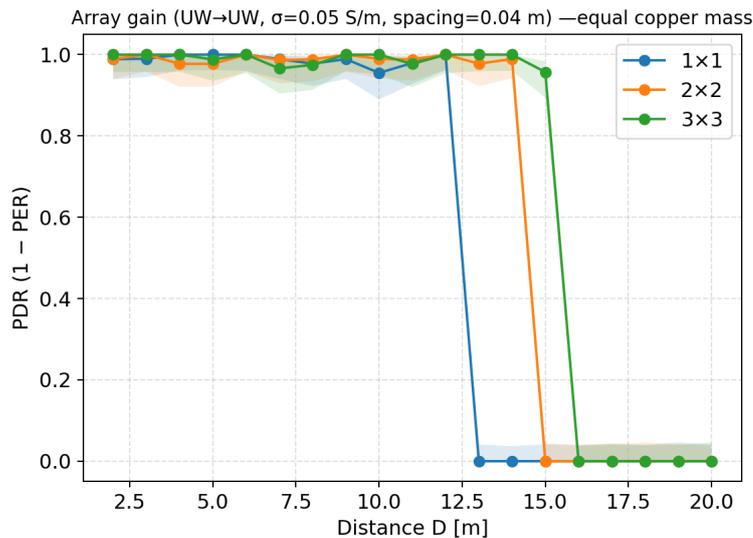


Figure 4.13: Array gain for UW→UW in freshwater: PDR vs distance for different numbers of coils per node with equal total copper mass ($\text{auto_scale_R}=1$, effective series resistance scaled with N).

ing the above-water node height between 0.2 and 2 m has only a modest effect on coverage. The PDR heatmaps exhibit an almost vertical good-PDR band, confirming that distance and conductivity dominate over small depth variations around the surface. This suggests that practical deployments can tolerate some motion of the surface node (e.g., waves and buoy oscillations) without substantial loss in range.

Fourth, the array experiments in Section 4.4 demonstrated that using multiple coils per node can provide tangible range gains. In freshwater, moving from a single coil to three coaxial coils per node extends the range from about 12 m to 17 m in the “raw coil count” (more-copper) case, and to about 15 m when the total copper mass is kept approximately constant by increasing the series resistance with the number of coils. The PDR curves remain very steep, so the main effect of the arrays is to shift the coverage edge outwards rather than to smooth the transition. These results quantify, in realistic system-level simulations, the array gains often reported analytically in the MI literature [8, 9, 5, 7].

Finally, the parameter sweeps in Section 4.5 highlighted which hardware knobs are most effective. In freshwater the range is almost flat between 20 and 500 kHz, whereas in seawater it degrades quickly with increasing frequency, confirming the need for low-frequency operation in highly conductive media. Within the explored range, variations in coil quality factor and receiver noise figure only change the range by about one metre. Likewise, trad-

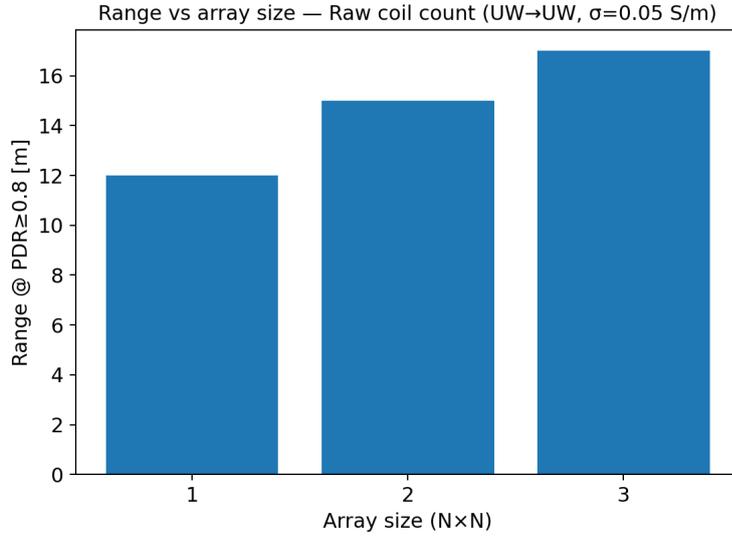


Figure 4.14: Range at $\text{PDR} \geq 0.8$ vs number of coils per node for the “raw coil count” configuration (UW→UW, $\sigma = 0.05$ S/m).

ing bandwidth for bit rate over $(B, R_b) \in \{2, 5, 8\} \text{ kHz} \times \{1, 2\} \text{ kbps}$ has a very minor impact on $D_{0.8}$. Overall, the simulations indicate that conductivity and coil geometry are the primary determinants of MI link range, while front–end noise and moderate changes in spectral efficiency act as second–order effects.

Beyond the numerical values, an important outcome of this work is the structure of the simulation framework itself. By exposing all physical and system parameters (coil geometry, conductivity, carrier frequency, noise figure, bandwidth, bit rate, number of coils and array mode) as Tcl options, and by enforcing a uniform CSV logging format across all experiments, the MI extensions developed in this thesis can be reused as a *design tool* rather than as a one–off scenario. New water types, node geometries or modem configurations can be explored by modifying only a few configuration lines, while the core propagation and PHY logic remain identical. This separation between a stable physics–based core and flexible high–level sweeps is intended to make MI studies in DESERT both rigorous and easy to extend.

These findings will be used in the Conclusions chapter to derive practical design guidelines for underwater MI links and to discuss the conditions under which MI communication is competitive with more traditional acoustic or RF underwater technologies [1, 2, 3, 4].

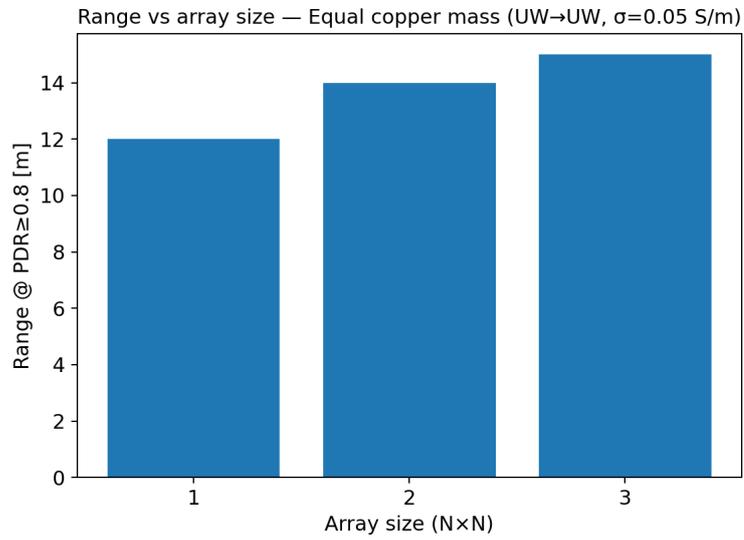


Figure 4.15: Range at $PDR \geq 0.8$ vs number of coils per node for the equal-copper-mass configuration.

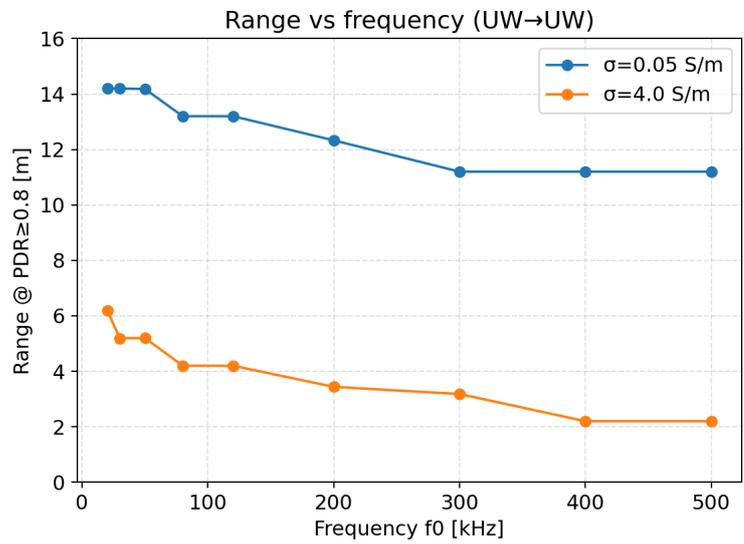


Figure 4.16: Range at $PDR \geq 0.8$ vs carrier frequency f_0 for the UW→UW link in freshwater and seawater.

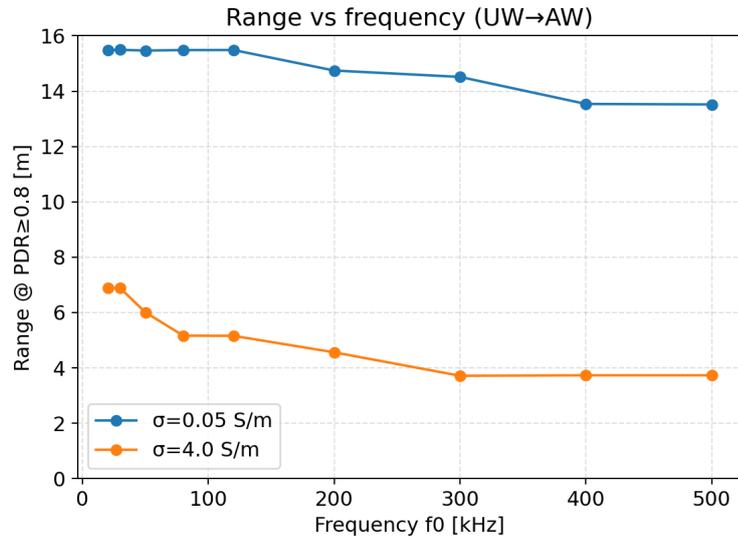


Figure 4.17: Range at PDR ≥ 0.8 vs f_0 for the UW→AW link in freshwater and seawater.

Table 4.2: Range $D_{0.8}$ vs quality factor Q for the UW→UW link in freshwater (PDR ≥ 0.8 , $f_0 = 200$ kHz, $R_b = 1$ kbps).

Q	$D_{0.8}$ [m]
20	12.1
30	12.1
40	12.1
50	13.0
70	12.1
100	12.1
150	12.1

Table 4.3: Range $D_{0.8}$ vs receiver noise figure (NF) for the UW→UW link in freshwater (PDR ≥ 0.8 , $f_0 = 200$ kHz, $R_b = 1$ kbps).

NF [dB]	$D_{0.8}$ [m]
10	12.1
14	12.1
18	13.0
22	12.1
26	12.1

Table 4.4: Range $D_{0.8}$ (in metres) for different (B, R_b) pairs in freshwater UW \rightarrow UW operation ($f_0 = 200$ kHz, $Q = 50, P_t = 0$ dBm).

B [kHz]	R_b [bps]	$D_{0.8}$ [m]
2	1000	13.18
2	2000	13.20
5	1000	12.34
5	2000	12.20
8	1000	12.19
8	2000	12.20

5

Conclusions and Future Work

5.1 CONCLUSIONS

The starting point of this thesis was a concrete question: under which conditions can two coils in water communicate reliably using magneto–inductive (MI) coupling, and how far can such a link be pushed with realistic hardware? Addressing this required following the chain from Maxwell’s equations and near–field coupling, through an equivalent circuit and link budget, to packet–level performance in a network simulator. The work has followed that path step by step and has produced a set of models, modules and scripts that connect these levels in a way that can be reused and extended.

On the modelling side, the thesis started from the magnetic dipole approximation for coaxial coils and from the behaviour of electromagnetic fields in conductive media. The familiar $1/d^3$ decay of mutual inductance in the near field was combined with skin–depth theory to express how conductive freshwater and seawater introduce an additional exponential attenuation. For underwater–to–air links a simple but effective two–layer geometry was introduced: only the segment of the path that lies underwater suffers conductive loss, while the segment in air is essentially lossless over the short ranges considered. These ingredients were then embedded in a lumped RLC model, in which each node is represented by a resonant RLC loop and the near–field interaction is captured by a mutual inductance term. This provides a direct bridge between electromagnetic intuition and the link–budget expressions used later in the simulator.

The second step was to translate this physical picture into a form usable inside the DESERT Underwater framework. Two new modules were implemented for this purpose. The propagation module, `UwMiCouplingPropagation`, evaluates the total mutual inductance between the coils, including the possibility of simple one-dimensional arrays and the two-layer water-air geometry, and delivers received power to the rest of the stack. The physical-layer module, `UwMiPhy`, converts this received power into E_b/N_0 for uncoded BPSK in AWGN and from there into bit and packet error probabilities. It exposes the usual control parameters of a practical modem, such as transmit power, carrier frequency, quality factor, noise figure, bandwidth and bit rate, and logs compact per-packet statistics so that performance can be analysed offline. Together, these modules implement a physics-based MI channel and a configurable MI PHY within DESERT, while remaining compatible with its existing MAC and routing mechanisms.

A third element of the work was the construction of an evaluation framework around these modules. Rather than tuning individual scenarios by hand, the thesis developed an automated sweep pipeline based on Tcl for scenario description, Bash for parameter sweeps and Python for post-processing. Distances, conductivities, depths, numbers of coils, frequencies, quality factors, noise figures and rates can all be swept systematically, with multiple random seeds per configuration. The Python scripts aggregate the results, compute Wilson confidence intervals for the packet delivery ratio and interpolate range metrics from the PDR curves, producing the tables and figures in Chapter 4. All of this is driven by CSV files, which makes the experiments reproducible and straightforward to extend: changing a coil radius, a water type or a modem parameter does not require altering the core logic of the simulations.

Within this framework, several conclusions emerge from the numerical results. A first consistency check is that the MI PHY behaves as expected in a purely AWGN setting. When the link budget inside `UwMiPhy` is flattened into E_b/N_0 , the simulated PER curves follow the analytical expression for uncoded BPSK over AWGN almost exactly, for both freshwater and seawater. The waterfall around $E_b/N_0 \approx 16\text{--}18$ dB matches the theoretical prediction for the packet sizes used in the thesis. This validates the link between received power, E_b/N_0 and PER, and ensures that any differences between scenarios can be attributed to the MI channel and not to artefacts of the PHY implementation.

The single-coil experiments then quantify what a small MI modem can achieve in practice. For the reference configuration used throughout the work, with circular coils of radius 0.12 m, about forty turns, series resistance around $4\ \Omega$, transmit power of 0 dBm and bit rate of 1 kbps, the achievable distance at high reliability is on the order of 12 m for underwater-

to–underwater operation in freshwater and about 14–15 m when one node is brought just above the surface. In seawater, where the skin depth is much shorter, the range shrinks to roughly 3–4.5 m under the same conditions. The PDR waterfalls, the median E_b/N_0 and the received–power curves are consistent with each other and with the theoretical model: the PDR drops when the received signal approaches the thermal noise floor, the median E_b/N_0 falls below the BPSK threshold, and the power decay follows a $1/d^6$ law in the near field with an additional exponential roll–off in seawater. This gives a concrete picture of what can be expected from small, low–power MI devices in realistic water conditions.

A further set of experiments isolates the role of the water–air interface and of node depth. Moving the receiver from just below the surface to up to two metres above it consistently improves the range by a modest but measurable amount, because part of the path is moved into air where conductive loss is negligible. However, once the receiver is some tens of centimetres above the surface, further increasing its height has only a limited effect compared with changing the horizontal distance or the conductivity. The heatmaps of PDR as a function of distance and receiver height show an almost vertical transition between good and poor reliability, which confirms that the link is primarily constrained by horizontal separation and medium properties. From an engineering point of view, this means that underwater–to–air MI links of the type considered here are relatively insensitive to wave–induced vertical motions of a surface buoy, provided that the separation and coil parameters are chosen appropriately.

The thesis also explores multi–coil arrays as a way to extend range. In the simple coaxial array model adopted, where identical coils are stacked along the link axis, increasing the number of coils per node proves to be a very effective design parameter. When the series resistance per node is kept equal to that of a single coil, adding more coils increases the effective magnetic moment and extends the freshwater underwater–to–underwater range from roughly 12 m to about 17 m when three coils are used at each end. Even when the total copper mass is kept approximately constant by scaling the resistance with the number of coils, the range still improves to around 14–15 m. In all cases the PDR curves remain steep, so the main effect of the array is to shift the coverage edge outwards rather than to smooth the transition. These results, obtained in a packet–level simulator, support and quantify the array gains that are often discussed analytically in MI studies.

Finally, parameter sweeps in carrier frequency, quality factor, noise figure, bandwidth and bit rate clarify which hardware parameters deserve most attention. In freshwater, the range remains almost flat over a wide frequency span between a few tens and a few hundreds of

kilohertz, indicating that once resonance is achieved the link is dominated by near-field geometry and coil size rather than by the precise choice of carrier. In seawater the situation is different: range decreases steadily as frequency increases, because the skin depth becomes comparable with or smaller than the link distance. In both media, varying the coil quality factor and the receiver noise figure within reasonable engineering ranges changes the range only by about a metre, which is comparable with the resolution of the discrete distance grid. Similarly, trading bandwidth for data rate within the tested (B, R_b) pairs has only a minor effect, especially once the resonant bandwidth of the coil limits the effective noise bandwidth. Taken together, these observations suggest a clear hierarchy of design levers: conductivity and coil geometry set the fundamental limits; multi-coil arrays are a useful way to push those limits; carrier frequency matters mainly in highly conductive water; and front-end details such as Q , noise figure and moderate rate-bandwidth choices act as second-order refinements.

All of these conclusions are subject to the deliberate simplifications of the model. The coils are assumed coaxial and perfectly aligned; the water is homogeneous within each layer and the water-air interface is flat; the PHY uses uncoded BPSK in AWGN and does not include interference or impulsive noise; and only static single-hop links are considered. Within this scope, however, the model passes several consistency checks and reproduces the order of magnitude of ranges reported in experimental MI prototypes with comparable parameters. The resulting framework is therefore not a complete description of all underwater effects, but a controlled and quantitatively meaningful approximation for short-range MI links with small coils and low to moderate data rates.

In summary, the thesis has shown that MI communication is a realistic option for short-range underwater links in low-conductivity water and, to a more limited extent, in seawater, provided that distances are kept in the range of a few to a few tens of metres and that coil design is chosen carefully. It has also demonstrated that this behaviour can be captured by a compact model integrated into DESERT Underwater and supported by an automated sweep framework. Beyond the specific numerical results, the main outcome is a bridge between near-field physics and network simulation for MI links, and a set of software artefacts that can be used as a starting point for further studies.

5.2 FUTURE WORK

The framework developed in this thesis was intentionally kept simple enough to be implemented and exercised within the time and resources available, yet modular enough to be

extended in several directions. The most natural extension is experimental validation. Building a small MI testbed with coils and front-end electronics similar to those assumed in the model, and performing controlled measurements in a tank or pool, would allow the propagation and noise parameters to be tuned against reality. Such an effort would either confirm the adequacy of the present abstractions or reveal systematic discrepancies due, for instance, to complex boundary conditions, manufacturing tolerances or unmodelled noise sources. In either case, the DESERT modules developed here would provide a convenient baseline.

On the propagation side, there is room to relax some of the geometric assumptions. Allowing arbitrary three-dimensional coil orientations, lateral offsets and random misalignment would make it possible to study scenarios in which the nodes are mounted on mobile robots, divers or drifting platforms, where perfect alignment cannot be guaranteed. Incorporating more detailed models of the water-air interface, seabed reflections or nearby structures would bring the simulations closer to coastal and harbour environments. These extensions would increase complexity, but they would naturally fit into the existing `UwMiCouplingPropagation` module.

At the physical-layer level, the current abstraction from received power to E_b/N_0 for uncoded BPSK can be reused for more advanced modems. Replacing the analytical PER curve with those corresponding to coded schemes or higher-order constellations, or adding simple baseband models where necessary, would allow the same MI channel to be evaluated under stronger coding and different waveform choices. This would help quantify how much additional range or reliability can be gained from coding and diversity techniques in the regimes identified here.

The existence of a configurable MI channel and PHY inside DESERT also opens the door to MAC and network-level studies. The present work has focused on single-hop links under controlled traffic, but the same modules can be used in multi-node topologies to investigate random access, scheduled MAC protocols, routing, topology control and mobility. In particular, hybrid scenarios where MI is combined with acoustic communication—for example, using MI for short-range, low-latency control and acoustic links for long-range backhaul—appear promising and could be explored without major changes to the existing code.

Finally, a broader line of future work is to place MI more explicitly within the landscape of underwater communication technologies. By combining the extensions developed here with DESERT's existing acoustic and RF models, it would be possible to perform systematic, like-for-like comparisons of range, energy per bit, latency and robustness across different media and deployment conditions. Such comparisons could help identify operating regimes where

MI is clearly advantageous, where it is competitive, and where it is best used as a complementary technology rather than as a replacement.

Underwater magneto–inductive communication will not make acoustic or optical systems obsolete, but it offers a distinct and useful region of the design space. By tying together a physics–based channel model, a configurable PHY and an automated simulation framework inside DESERT Underwater, this thesis has provided a compact environment for exploring that region. The intention is that the models and tools left behind will make it straightforward for others to adapt the framework to their own coils, water types and applications, and to continue refining our understanding of when, and how, MI links are worth deploying in real underwater networks.



MI Simulation Framework and Reproducibility

This appendix provides a brief practical guide to the software artefacts developed in this thesis and to the procedure used to generate the numerical results of Chapter 4. It is not a full DESERT Underwater tutorial, but rather a focused description of how the magneto-inductive (MI) extensions are organised and how they can be reused as a small design framework.

A.1 CODE COMPONENTS AND ROLES

The MI simulation framework built in this work consists of four main blocks:

1. MI propagation module (**UwMiCouplingPropagation**). A C++ class derived from DESERT's `MPropagation` that implements the near-field MI coupling law, coil arrays and the two-layer underwater-to-air loss model described in Chapter 2. It exposes Tcl parameters for coil geometry, water conductivity, array size, resistance scaling and the optional conductive-loss term.
2. MI physical layer (**UwMiPhy**). A C++ class derived from `MPhy_Bpsk` that converts the received power from the propagation model into E_b/N_0 , BER and PER for uncoded BPSK over AWGN, as detailed in Chapter 3. It exposes Tcl parameters for transmit power, bit rate, centre frequency, quality factor, effective noise bandwidth, noise fig-

ure and equivalent noise temperature, and logs compact per–packet metrics to text files.

3. Tcl scenario (**test_uwmi_net.tcl**). A Tcl script that assembles the end–to–end DESERT stack (CBR–UDP–IP–CSMA/ALOHA–MLL–UwMiPhy–UwMiCouplingPropagation) for a two–node link, attaches position objects, constructs the rectangular spectral mask and passes all MI parameters as command–line options. This script is the main entry point for single simulation runs.
4. Sweep and post–processing scripts (**sweep_runs_v2.sh**, **make_core_plots.py**). A Bash wrapper that repeatedly calls the Tcl scenario over grids of distances and parameters (conductivity, frequency, number of coils, quality factor, noise figure, bandwidth and bit rate) and a Python script that aggregates the resulting CSV logs and produces the figures and tables of Chapter 4. These scripts implement the automation pipeline summarised in Figure 3.2.

The C++ modules are compiled as part of the DESERT Underwater tree, while the Tcl and Bash/Python scripts can be kept in a separate directory, as long as the paths to the ns binary and to the Tcl scenario are correctly configured in `sweep_runs_v2.sh`.

A.2 RUNNING A SINGLE MI SIMULATION

A single two–node MI simulation can be run directly through the Tcl scenario. The script accepts a number of command–line flags, which are grouped in Chapter 3, Section 3.2.3. Typical examples include:

- geometry and medium: `-distance`, `-z1`, `-z2`, `-uw2aw`, `-sigma`;
- RF and noise: `-f_khz`, `-B`, `-Rb`, `-Q`, `-NF_dB`, `-Tnoise_K`, `-txpwr_dbm`;
- coil and array parameters: `-Nt`, `-Nr`, `-at`, `-ar`, `-Rt`, `-Rr`, `-kappa`, `-Nt_coils`, `-Nr_coils`, `-st`, `-sr`, `-auto_scale_R`;
- simulation control: `-pktsize`, `-period`, `-maxpkts`, `-app_start`, `-app_stop`, `-sim_stop`, `-seed`.

A typical invocation for a $D = 10$ m UW→UW link in fresh water with default hardware parameters is:

```

ns test_uwmi_net.tcl \
  -distance 10.0 -z1 -5.0 -z2 -5.0 -uw2aw 0 \
  -sigma 0.05 \
  -f_khz 200 -B 5000 -Rb 1000 -Q 50 \
  -Nt 40 -Nr 40 -at 0.12 -ar 0.12 \
  -Rt 4.0 -Rr 4.0 -kappa 0.7 \
  -txpwr_dbm 0.0 \
  -pktsize 128 -period 0.5 -maxpkts 2000 \
  -app_start 10.0 -app_stop 1010.0 -sim_stop 1020.0 \
  -seed 1

```

At the end of the run, the script writes a per-run summary line to a CSV file (recording PDR, geometry and parameters) and a more detailed PHY log with one UWMI_METRIC line per packet, as described in Section 3.5.

A.3 AUTOMATED SWEEPS WITH `sweep_runs_v2.sh`

For the experiments in Chapter 4, the Tcl scenario is wrapped by a Bash script `sweep_runs_v2.sh`. The script sets common options (carrier frequency, bandwidth, bit rate, packet size, number of packets per run, number of seeds, output directory) and then defines helper functions such as:

- `run_ns`: a thin wrapper around the ns invocation, passing through the MI parameters and logging directories;
- `run_core_grids`: sweeps the distance for the four combinations of mode (UW→UW and UW→AW) and conductivity (fresh and salt water) used in Section 4.2;
- `sweep_uwa_depth_family`: sweeps the receiver height above the surface, as in Section 4.3;
- `sweep_arrays`: sweeps the number of coils per node and the `auto_scale_R` flag, as in Section 4.4;
- `sweep_Q`, `sweep_NF` and `sweep_B_Rb_grid`: perform the ablation studies on quality factor, noise figure and (B, R_b) pairs reported in Section 4.5.

The bottom of the script implements a simple command-line interface that dispatches to these helpers depending on the first argument (e.g., `core`, `freq`, `arrays`). Inspecting `sweep_runs_v2.sh` shows the available modes and the corresponding calls.

A typical usage pattern is:

```
# Run the core distance sweeps in all four (mode, sigma) combinations
./sweep_runs_v2.sh core

# Sweep carrier frequency and regenerate the range vs frequency plots
./sweep_runs_v2.sh freq
```

Each call produces a set of ns-2 runs, organised in subdirectories by mode and conductivity, and appends results to a master CSV file such as `uwmi_results_summary_all.csv`.

A.4 POST-PROCESSING AND FIGURE GENERATION

The Python script `make_core_plots.py` parses the CSV summaries and per-packet logs produced by the sweeps and generates the figures and tables in Chapter 4. Its main steps are:

1. Aggregation of PDR. Group runs by $(D, \sigma, \text{mode}, f_0, R_b, B)$ and accumulate total packets sent and received over seeds, yielding empirical PDR and Wilson confidence intervals, as detailed in Section 3.7.
2. Extraction of PHY metrics. Scan all `UWMI_METRIC` lines in the logs to recover per-packet received power, noise power and E_b/N_0 , and compute median and interquartile ranges as a function of distance.
3. Range and design tables. Apply the range extractor `range_from_P()` to obtain D_τ for chosen PDR thresholds, and assemble the range-versus-parameter tables (e.g., range vs frequency, vs quality factor, vs noise figure, and the rate-range table).
4. Plotting. Generate the PDR vs distance curves with shaded confidence bands, the E_b/N_0 and received-power vs distance plots with the $1/d^6$ reference slope and thermal-noise floor, the depth heatmaps, the array curves and the parameter sensitivity plots.

After the sweeps have completed, regenerating the core figures typically amounts to running:

```
python3 make_core_plots.py
```

from the directory where the CSV and log files have been collected. The script writes all intermediate tables (e.g., range vs frequency, range vs Q , range vs NF, rate-range frontier) as CSV files in addition to the PDF or PNG figures, so that the numerical results can easily be reused.

A.5 ADAPTING THE FRAMEWORK TO OTHER MI DESIGNS

A key goal of this work is that the MI extensions should be usable as a small framework, rather than a single hard-coded scenario. Adapting the simulations to a different MI modem or deployment typically requires changing only:

- *Tcl parameters* for coil geometry and hardware (e.g., `-at`, `-ar`, `-Nt`, `-Nr`, `-Rt`, `-Rr`, `-Q`, `-NF_dB`, `-txpwr_dbm`);
- *sweep ranges* in `sweep_runs_v2.sh` (e.g., the list of distances, frequencies or array sizes);
- optionally, *plot labels and axis limits* in `make_core_plots.py`.

The underlying MI physics (mutual inductance, skin depth, conductive loss) and the PER model (uncoded BPSK over AWGN) are encapsulated inside `UwMiCouplingPropagation` and `UwMiPhy` and need not be modified. In this sense, the framework can be viewed as a reusable building block: once compiled into DESERT, it can support future studies of coded PHYs, alternative MAC protocols, multi-hop MI networks or hybrid acoustic-MI architectures by plugging into different stacks and sweep scripts.

References

- [1] I. F. Akyildiz, D. Pompili, and T. Melodia, “Underwater acoustic sensor networks: research challenges,” *Ad Hoc Networks*, vol. 3, no. 3, pp. 257–279, 2005.
- [2] J. Heidemann, M. Stojanovic, and M. Zorzi, “Underwater sensor networks: applications, advances and challenges,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 158–175, 2012.
- [3] A. Pal, F. Campagnaro, K. Ashraf, M. R. Rahman, A. Ashok, and H. Guo, “Communication for underwater sensor networks: A comprehensive summary,” *ACM Transactions on Sensor Networks*, vol. 19, 2023, article 22, 44 pages.
- [4] Y. Lou and N. Ahmed, *Underwater Communications and Networks*. Cham: Springer, 2022.
- [5] M. C. Domingo, “Magnetic induction for underwater wireless communication networks,” *IEEE Transactions on Antennas and Propagation*, vol. 60, no. 6, pp. 2929–2939, 2012.
- [6] Z. Sun and I. F. Akyildiz, “Realizing underwater communication through magnetic induction,” *IEEE Communications Magazine*, vol. 53, no. 11, 2015.
- [7] Y. Li, X. Liang, W. Guo, Z. Liu, B. M. Chen, M. Chen, and Z. Sun, “Underwater magnetic induction communications: A survey,” *IEEE Communications Surveys and Tutorials*, vol. 21, no. 4, pp. 3416–3441, 2019.
- [8] Z. Sun and I. F. Akyildiz, “Magnetic induction communications for wireless underground sensor networks,” *IEEE Transactions on Antennas and Propagation*, vol. 58, no. 7, pp. 2426–2435, 2010.
- [9] B. Gulbahar and O. B. Akan, “A communication theoretical modeling and analysis of underwater magneto-inductive wireless channels,” *IEEE Transactions on Wireless Communications*, vol. 11, no. 9, pp. 3326–3334, 2012.

- [10] I. Calabrese, R. Masiero, P. Casari, L. Vangelista, and M. Zorzi, “Embedded systems for prototyping underwater acoustic networks: The DESERT Underwater libraries on board the pandaboard and netdcu,” in *Proceedings of MTS/IEEE OCEANS*, Yeosu, Korea, 2012.
- [11] N. Ahmed, J. Hoyt, A. Radchenko, D. Pommerenke, and Y. R. Zheng, “A multi-coil magneto-inductive transceiver for low-cost wireless sensor networks,” in *Proceedings of the 2014 Underwater Communications and Networking Conference (UComms)*, 2014.
- [12] N. Ahmed, G. Qiao, Y. R. Zheng, and D. J. Pommerenke, “Design and implementation of medium access control protocol for magneto-inductive wireless sensor networks using low power sensor nodes,” *IEEE Journal of Oceanic Engineering*, 2023, early access.
- [13] M. Aman, G. Qiao, Z. Shang, Z. U. Khan, M. S. Khan, I. Ullah, and Wazir-Ur-Rahman, “Realization of rssi based, three major components (hx, hy, hz) of magnetic flux created around the mi-td coil,” in *Proceedings of the 2023 IEEE International Conference on Electrical, Computer and Energy Engineering (ICEACE)*, 2023.
- [14] C. A. Balanis, *Advanced Engineering Electromagnetics*, 2nd ed. Wiley, 2012.
- [15] J. G. Proakis and M. Salehi, *Digital Communications*, 5th ed. McGraw-Hill, 2008.

Acknowledgments

I begin by apologising to my family for the time I have not been with you, following a dream that has often felt more like destiny than choice. Every missed moment with you is written quietly between these lines.

To my supervisor and professors, thank you for your guidance, patience and rigour; for setting a standard of excellence that has shaped the way I think and work; and for trusting me with opportunities I once saw only from a distance.

To my extended family, my friends, my colleagues and all those who crossed my path on this journey — some close, some far, some almost strangers — thank you. Your presence, words and kindness have meant more than you know. To Xheina, thank you for sharing this journey with me: we cried, we laughed, and we managed to be each other's hand in the middle of simply trying to survive.

To my family: you know the story behind this work, even when it is not on the page. You are the reason this dream was born and the reason it continued. This is for you.

To my sister Ana, my fearless defender and biggest cheerleader, who always believed I was meant for more and never let me forget it. You have been my courage when mine was low.

I also want to thank the part of myself I often overlook: Vanesa, who has rarely taken a real day off; who used holidays to sit exams; who travelled, studied and worked across two countries when it would have been easier to stop. Thank you for being a fighter, for your passion, and for carrying us here.

Finally, to my family once more: I know that distance has its cost. I hope, one day, to repay in presence the years I have spent far from you, and to show that the path we chose together was worth it.