

SCHOOL OF ELECTRICAL ENGINEERING AND TELECOMMUNICATIONS

Joint Trajectory and SDMA Beamforming for Secure Full-Duplex UAV Networks Against Active Eavesdropping

by

Yan Liang

Thesis submitted as a requirement for the degree
Master of Engineering (Telecommunications Engineering)

Supervisor: A/Prof. Derrick Wing Kwan Ng

Submitted: October 25, 2026

Abstract

Unmanned aerial vehicles (UAVs) offer controllable mobility and line-of-sight (LoS) channels, and full-duplex (FD) operation with space-division multiple access lets a single UAV serve multiple downlink and uplink users in the same time–frequency resource. The same LoS propagation, however, exposes the aerial link to eavesdropping—and an FD-capable eavesdropper can jam and intercept simultaneously, a strictly stronger threat than the passive adversary assumed in most prior work.

This thesis jointly designs the UAV trajectory, SDMA beamforming, and artificial-noise (AN) allocation—including when AN should be allocated at all—to minimize the UAV's total propulsion and communication energy under per-user rate, per-slot secrecy, power, self-interference, and kinematic constraints, against a ground eavesdropper that may itself be FD and active. The non-convex problem is solved by a four-block alternating optimization combining semidefinite relaxation and successive convex approximation; the objective sequence is monotonically non-increasing and convergent, with feasibility guaranteed at each cycle and limit points that are block-coordinatewise stationary for the relaxed problem under standard surrogate conditions.

Analyzing the secrecy constraints in closed form, we derive an offline-checkable condition under which AN is provably redundant: it reduces to comparing the eavesdropper's own residual self-interference with the uplink power it intercepts—a self-defeating property of FD active eavesdropping that we isolate in closed form. In the complementary regime where AN becomes a mandatory friendly-jamming floor, the required power grows quadratically with the mission-enforced UAV–eavesdropper distance, rendering this regime energetically unreachable under practical power budgets and per-slot secrecy guarantees. Numerically, the joint design reduces total energy by 9.2% over a beamforming-only baseline in a representative secrecy-constrained scenario, and we show how the optimized trajectory is progressively reshaped by transmit power, per-user demand, and eavesdropper proximity.

Acknowledgements

There is a line of classical Chinese poetry that I have kept with me: *Having known the vastness of the world, I cherish still the green of grass and trees.*

Two years of postgraduate study at UNSW have broadened my academic horizons, showing me more facets of telecommunications than I knew existed, and more possibilities for where the field might go. What I hope to carry forward, in my research and in my life, is the feeling I had in the first moment I walked through the University's gates: that sense of anticipation, and that curiosity.

My deepest thanks go to my supervisor, Associate Professor Derrick Wing Kwan Ng, who has been both a mentor and a friend. Having him as my supervisor was the most fortunate thing to happen to me during my master's degree. Academically, the depth of his knowledge and technical insight gave me both a tremendous sense of security and a clear sense of direction; personally, his encouragement and support meant that I was able to enjoy every stage of writing this thesis rather than merely endure it. From the shape of an idea to the mathematics behind it, his help has been beyond measure. Most importantly, he taught me that good research is a rigorous undertaking that demands sustained and substantial time — a lesson I will carry with me for the rest of my life. I am also grateful to the School of Electrical Engineering and Telecommunications at UNSW for the resources and support it gave me throughout this research.

To my family: thank you for the encouragement and support that gave me the courage to look at this world once more. To my friends: thank you for carrying the heaviest moments with me, and for raising a glass with me at every small victory. To my partner: thank you for always believing I was more resilient and more capable than I believed myself to be — and then, somehow, I became that person. Thank you for each dinner with me. And thank you to myself — for being willing to begin again. Onward to the next journey, without fear.

Contents

Abstract **1**

Acknowledgements **2**

List of Figures **5**

List of Tables **6**

Abbreviations **8**

1 Introduction **9**

1.1 Motivation 9

1.2 Problem Statement 10

2 Background **11**

2.1 UAV-Enabled Wireless Communication 11

2.2 Full-Duplex Communication and Self-Interference 12

2.3 Physical-Layer Security Fundamentals 12

2.4 Multiuser Beamforming and SDMA 13

2.5 Optimization Tools for Non-Convex Trajectory and Beamforming Design . 14

3 Literature Review **16**

3.1 UAV mobility, propulsion energy, and resource allocation 16

3.2 Multiuser beamforming and full-duplex aerial communications 17

3.3 Trajectory optimisation and physical-layer security for UAVs 18

3.4 Energy-aware secure transmission and covert alternatives 19

3.5 Passive, active, and full-duplex eavesdropper models 20
3.6 Robust beamforming, artificial noise, and optimisation methods 21
3.7 Synthesis and research gap 22

4 System Model **24**
4.1 Channel model 24
4.2 Signal model and SINRs 25
4.3 Problem Formulation 27

5 Proposed Algorithm **28**
5.1 Rate DC reformulation and the common SCA surrogate 28
5.2 Block 1a: information beamforming 30
5.3 Block 1b: AN covariance update 31
5.4 Block 2: uplink combiner update 31
5.5 Block 3: trajectory and kinematics 32

6 Convergence Analysis **35**

7 Secrecy Analysis **37**
7.1 The friendly-jamming threshold c_u 37
7.2 The AN-redundancy condition 38
7.3 Verification in the studied configurations 38
7.4 The mandatory regime: reachability and a power scaling law 39
7.5 Numerical caveat 40

8 Results **41**
8.1 Simulation Setup 41
8.2 How the trajectory is shaped 41
8.3 Benchmark: value of joint optimization 44
8.4 Artificial noise in the numerical results 47

9 Conclusion **48**

9.1 Future Work 49

List of Figures

4.1 System model: a full-duplex UAV with N_t transmit and N_r receive antennas flies from $\mathbf{q}_I$ to $\mathbf{q}_F$ at altitude H, simultaneously serving K downlink and U uplink ground users over LoS channels ($\mathbf{h}_k$, $\mathbf{g}_u$), while a full-duplex active eavesdropper overhears both the downlink ($\mathbf{h}_E$) and the uplink ($g_{E,u}$). . . . 25

5.1 AO framework with SDR-based transmit design, closed-form uplink combining, and SCA-based trajectory optimization. The objective is non-increasing over iterations and the loop terminates when the relative improvement falls below the tolerance ϵ. 29

8.1 Trajectory shaping under increasing physical factors: (a) high power → straight; (b) normal power → bends to users; (c) high-demand user → demand-weighted detour; (d) with Eve → compromise between service and secrecy. Color encodes UAV speed. 43

8.2 Trajectory comparison of the three schemes (speed-colored). 44

8.3 Final energy breakdown. Left: total energy (kJ); right: transmit energy (J). 45

8.4 AO convergence: total energy decreases monotonically and converges within a few iterations for all three schemes. 46

8.5 Final average DL/UL rates. All schemes meet the QoS floors: the DL rates settle near $R_{\min}^{\mathrm{DL}}$, while the UL rates lie well above $R_{\min}^{\mathrm{UL}}$. 46

List of Tables

3.1 Representative studies and their relationship to the present formulation. . . 23

8.1 Baseline parameters. 41

8.2 Total energy under the safety scenario ($R_{\min}^{\mathrm{DL}} = 0.3$, Eve near path, $\Gamma_E^{\mathrm{DL}} = 0.2$). 44

Abbreviations

AN	Artificial Noise
AO	Alternating Optimization
BSUM	Block Successive Upper-bound Minimization
CVX	Convex optimization solver (MATLAB toolbox)
DC	Difference-of-Convex
DL	Downlink
FD	Full-Duplex
IoT	Internet of Things
KKT	Karush–Kuhn–Tucker
LoS	Line-of-Sight
MIMO	Multiple-Input Multiple-Output
MISO	Multiple-Input Single-Output
MMSE	Minimum Mean Square Error
NOMA	Non-Orthogonal Multiple Access
PA	Power Amplifier
QoS	Quality of Service
RIS	Reconfigurable Intelligent Surface
RSI	Residual Self-Interference
SCA	Successive Convex Approximation
SDMA	Space-Division Multiple Access
SDP	Semidefinite Program
SDR	Semidefinite Relaxation
SI	Self-Interference
SINR	Signal-to-Interference-plus-Noise Ratio
UAV	Unmanned Aerial Vehicle
UL	Uplink
5G	Fifth-Generation (wireless network)

Chapter 1

Introduction

1.1 Motivation

Unmanned aerial vehicles (UAVs) have emerged as a versatile platform for wireless communication, owing to their low deployment cost, high mobility, and controllable three-dimensional positioning [1]. Relative to terrestrial base stations, a UAV can be dispatched on demand, repositioned to track traffic and channel conditions, and, because it typically operates well above ground clutter, is far more likely to maintain a line-of-sight (LoS) link to the users it serves. These properties have made UAV-assisted communication a central theme in the design of fifth-generation (5G) and beyond wireless networks, spanning applications from rapid coverage restoration after disasters to aerial relaying and data collection for the Internet of Things [2].

A UAV's usefulness as a communication platform is, however, constrained by three coupled factors. First, its onboard battery imposes a hard energy budget that is dominated by propulsion rather than communication, so that trajectory design has a first-order effect on mission endurance [3]. Second, a single low-altitude platform serving a cluster of ground users must share limited spectral resources among them; space-division multiple access (SDMA), in which spatial beamforming rather than orthogonal time or frequency resources separates users, together with full-duplex (FD) operation, in which the UAV transmits and receives in the same time-frequency resource, are two complementary ways to raise the spectral efficiency of the aerial link without proportionally consuming more of the UAV's finite time budget [4]. Third, the same LoS propagation that makes UAV channels attractive for legitimate users also makes them highly vulnerable to interception: a ground-based eavesdropper is statistically far more likely to enjoy a strong, unobstructed channel to an aerial transmitter than to a terrestrial one shielded by buildings and terrain. Physical-layer security, which exploits the multi-antenna and trajectory degrees of freedom of the UAV itself rather than upper-layer cryptography, is therefore a natural complement

to UAV communication design [5, 6].

These three factors interact in ways that cannot be optimized separately. Flying to minimize propulsion energy pulls the UAV away from users and toward a straight, high-speed path; flying to maximize legitimate throughput pulls it toward users; and flying to preserve secrecy pushes it away from any eavesdropper that happens to lie between the UAV and its users. Full-duplex operation compounds the difficulty: it requires the UAV to suppress its own residual self-interference (RSI) while simultaneously serving downlink (DL) and uplink (UL) users, and, critically, it hands the same capability to a sufficiently sophisticated adversary. An *active* eavesdropper equipped with full-duplex hardware can jam the legitimate uplink while simultaneously listening to the downlink, or vice versa, a threat that is qualitatively different from the passive, receive-only eavesdroppers assumed in most of the UAV secrecy literature [7, 8]. Designing the UAV's trajectory, its multiuser SDMA beamformers, and its artificial-noise (AN) allocation jointly, under an energy budget, and against an active full-duplex adversary, is therefore a substantially richer problem than any one of trajectory optimization, secure beamforming, or full-duplex resource allocation considered in isolation.

1.2 Problem Statement

This thesis considers a single full-duplex UAV that simultaneously serves K downlink users and U uplink users while a ground eavesdropper attempts to intercept either link. The UAV flies at fixed altitude over a finite mission horizon, discretized into N time slots, and is equipped with multiple transmit and receive antennas so that it can apply SDMA beamforming to the downlink users, allocate artificial noise to degrade the eavesdropper's channel, and combine its uplink receive antennas to suppress residual self-interference and multiuser interference. The design objective is to minimize the UAV's total energy consumption – propulsion plus communication – subject to per-user quality-of-service (QoS) constraints on both links, secrecy constraints that cap the information leaked to the eavesdropper on both links, the UAV's kinematic limits, and its transmit power budget. The resulting problem is non-convex in the trajectory, the downlink covariances, the artificial-noise covariance, and the uplink combiners jointly, and is further complicated by the possibility that the eavesdropper itself operates in full-duplex mode, which changes the structure of the uplink secrecy constraint from a simple upper bound on the UAV's transmit power toward Eve into a two-sided requirement that can, depending on the eavesdropper's self-interference cancellation capability, either forbid or mandate artificial-noise radiation in Eve's direction.

Chapter 2

Background

2.1 UAV-Enabled Wireless Communication

The defining feature of a UAV-enabled wireless link, relative to a terrestrial one, is that the air-to-ground channel is dominated by a line-of-sight path. Under this assumption, the channel between the UAV and a ground node at horizontal distance $d_{2\mathrm{D}}$ and UAV altitude H is commonly modeled as a free-space path-loss link with power gain $\beta_0 d^{-\alpha}$, where $d = \sqrt{H^2 + d_{2\mathrm{D}}^2}$ is the three-dimensional distance, β_0 is the channel power gain at a reference distance of one metre, and α is the path-loss exponent (with $\alpha = 2$ recovering the free-space case) [1, 3]. Because $\beta_0 d^{-\alpha}$ decreases monotonically with distance, the UAV's horizontal position directly controls the achievable rate to every ground node it serves, which is what makes trajectory design a communication-layer, and not merely a mobility-layer, decision.

The second defining feature of UAV communication is that the UAV's motion itself consumes energy, and this propulsion energy typically dominates the communication energy by two to three orders of magnitude. For a fixed-wing UAV flying at constant altitude, level flight, and negligible wind, the instantaneous propulsion power as a function of flight speed $v = \|\mathbf{v}\|$ is commonly approximated as

$$P(v) = c_1 v^3 + \frac{c_2}{v} + c_3, \tag{2.1}$$

where $c_1 v^3$ captures the parasitic (profile) drag power that grows with speed, c_2/v captures the induced drag power that grows as speed decreases, and c_3 collects speed-independent terms [3]. This power curve has a unique minimizer $v^\star = (c_2/3c_1)^{1/4}$, so that, absent any communication requirement, an energy-minimal UAV would simply cruise at $v^\star$. Communication requirements perturb the trajectory away from this unconstrained optimum by requiring the UAV to slow down or divert toward users, which is precisely the coupling

that joint trajectory-and-communication design must resolve.

2.2 Full-Duplex Communication and Self-Interference

A conventional (half-duplex) wireless node transmits and receives on orthogonal time or frequency resources, halving its effective spectral efficiency relative to an idealized simultaneous-transmit-and-receive system. Full-duplex (FD) radios remove this restriction by transmitting and receiving on the same time-frequency resource, at the cost of self-interference (SI): the node's own transmitted signal leaks into its receive chain at a power level that can exceed the desired received signal by many tens of decibels, since the transmit and receive antennas are co-located. Practical FD systems combine passive isolation, analog cancellation, and digital cancellation to suppress this SI to a residual level, commonly denoted the residual self-interference (RSI), that is treated as an additional noise-like term at the receiver.

For a UAV platform, the SI channel is aggravated by the platform's own vibration and by continuous relative motion, which makes the SI channel time-varying and complicates the channel-estimation step that underlies digital SI cancellation; dedicated pilot-aided digital cancellation schemes have been proposed to track this time-varying SI channel for FD UAV transceivers [9]. In the system model adopted in this thesis, the net effect of imperfect cancellation is captured by a residual self-interference term, scaled by a self-interference cancellation factor ϕ, that appears in the signal-to-interference-plus-noise ratio (SINR) of whichever link – uplink reception at the UAV, or downlink reception at a legitimate user – coexists with the UAV's simultaneous transmission. Crucially, the same physical mechanism applies to *any* full-duplex node, including an adversary: if the eavesdropper itself is a full-duplex device, it too suffers RSI when it jams and eavesdrops simultaneously, a property that Section 2.3 and the algorithmic chapters exploit rather than merely tolerate.

2.3 Physical-Layer Security Fundamentals

Physical-layer security seeks to guarantee confidentiality using the statistical and structural properties of the wireless channel itself, rather than (or in addition to) upper-layer cryptography. The foundational result is Wyner's wiretap channel [5], which shows that a positive secrecy rate – a rate at which a legitimate receiver can decode reliably while an eavesdropper with a degraded channel learns asymptotically nothing – is achievable whenever the eavesdropper's channel is statistically worse than the legitimate channel. In a UAV setting, however, both the legitimate and eavesdropping channels are often near-

LoS, so relying on inherent channel degradation is unreliable; instead, the transmitter must actively engineer a channel advantage.

Two complementary mechanisms are used for this purpose in this thesis. The first is spatial beamforming: with multiple transmit antennas, the UAV can shape its downlink signal to concentrate energy toward the legitimate user and place a null, or a low-gain direction, toward the eavesdropper, provided the number of transmit antennas exceeds the number of spatial directions that must be simultaneously served and nulled. The second is artificial noise (AN), introduced by Goel and Negi [6]: the transmitter deliberately radiates a noise-like signal, designed within the null space of the legitimate channel (or, more generally, so as to minimize its impact on the legitimate user while maximally degrading the eavesdropper), so that even if a beamforming null is unavailable or insufficient, the eavesdropper's received SINR can still be suppressed. In the uplink direction, an analogous but structurally different mechanism applies: because the UAV cannot beamform *away* from an uplink transmitter it needs to receive from, protecting an uplink transmission against an eavesdropper instead typically requires *friendly jamming*, in which the UAV (or another node) deliberately radiates interference toward the eavesdropper's receiver while the legitimate uplink receiver, whose location and combiner are known, remains largely unaffected. As later chapters show, whether AN is best used as a downlink null-space signal or as an uplink friendly-jamming floor depends on the sign of an explicit, offline-computable quantity that captures whether the eavesdropper's own receiver is already interference-limited.

Against a full-duplex active eavesdropper, this picture is further complicated in a way that is favourable to the legitimate system: an FD eavesdropper that jams while it listens necessarily contaminates its own receive chain with its own RSI, exactly as a legitimate FD node would. This self-defeating property, analyzed by Mukherjee and Swindlehurst for the general MIMO wiretap channel [7,8], means that an active eavesdropper's jamming power is not a free parameter it can increase without limit; it is fundamentally in tension with the eavesdropper's own ability to cancel its self-interference. This thesis specializes and quantifies that tension for the UAV geometry considered here.

2.4 Multiuser Beamforming and SDMA

When a single aerial transmitter must serve multiple ground users simultaneously and within the same frequency-time resource, some form of multiple access is required to separate their signals. Orthogonal schemes (time- or frequency-division multiple access) avoid inter-user interference by construction but under-utilize the spatial degrees of freedom offered by a multi-antenna UAV. Power-domain non-orthogonal multiple access (NOMA)

instead separates users by power level and successive interference cancellation. Space-division multiple access (SDMA), the scheme adopted in this thesis, separates users spatially: with N_t transmit antennas, the UAV forms a distinct beamforming vector (or, more generally, a transmit covariance matrix) for each downlink user, so that all users are served in the same time-frequency resource, with inter-user interference controlled through the spatial precoding rather than through orthogonal resource allocation [4]. The trade-offs among orthogonal, non-orthogonal, and spatial-domain access schemes, together with the resource-allocation machinery each of them requires, are surveyed comprehensively in [10]; the spatial-domain design adopted here inherits both the interference-control burden and the degrees of freedom described there.

Designing an SDMA transmit covariance $\mathbf{W}_k$ for user k that simultaneously maximizes user k's desired signal, minimizes leakage to co-scheduled users, respects a secrecy constraint against an eavesdropper, and remains within a total or per-antenna power budget is, in general, a non-convex rank-constrained problem, because the physically meaningful transmit signal corresponds to a rank-one covariance $\mathbf{W}_k = \mathbf{w}_k\mathbf{w}_k^{\mathsf{H}}$. Section 2.5 summarizes how this rank constraint is relaxed in practice.

2.5 Optimization Tools for Non-Convex Trajectory and Beamforming Design

The joint trajectory-SDMA-AN design problem outlined in Chapter 1 is non-convex in essentially every one of its variable groups: the achievable-rate expressions are non-concave fractional (SINR-based) functions of the beamforming covariances, the transmit covariances are subject to a non-convex rank-one constraint, and the propulsion-energy and channel-gain terms are non-convex functions of the trajectory. Three optimization techniques, used throughout the algorithmic chapters of this thesis, make the problem tractable.

Semidefinite relaxation (SDR). Beamforming problems posed in terms of a rank-one covariance $\mathbf{W}_k = \mathbf{w}_k\mathbf{w}_k^{\mathsf{H}}$ can be relaxed by dropping the rank-one constraint and optimizing instead over the convex cone of positive semidefinite matrices $\mathbf{W}_k \succeq 0$. The relaxed problem is a semidefinite program (SDP), solvable to global optimality by interior-point methods; Luo *et al.* survey the conditions under which the relaxed solution is either exactly rank-one or can be converted to a feasible rank-one solution via Gaussian randomization, and this SDR-plus-randomization pipeline underlies the downlink beamforming sub-problem solved in this thesis [11].

Successive convex approximation (SCA) and DC programming. Achievable-rate constraints of the form $\log_2(1 + S/I) \geq R_{\min}$ are non-convex because they involve a ratio of two functions of the optimization variables. Writing the rate as a difference of two concave functions, $R = \log_2(S + I) - \log_2(I)$, expresses the non-convexity in a difference-of-convex (DC) form; linearizing the concave, subtracted term around the current iterate produces a global concave lower bound on the true rate that is tight at the linearization point. Replacing the original rate constraint with this concave surrogate at each iteration, and re-linearizing after every update, is the essence of SCA, and is the mechanism used both for the SDMA/AN rate constraints and for the propulsion-energy and channel-gain terms in the trajectory sub-problem of this thesis.

Alternating optimization and block successive minimization. Because the joint problem is non-convex across all variable groups simultaneously but convex (after SDR/SCA) within each group given the others fixed, an alternating-optimization (AO) scheme – cyclically updating one block of variables while holding the rest fixed – is a natural solution strategy. Razaviyayn, Hong and Luo's block successive upper-bound minimization (BSUM) framework gives general conditions under which such block-coordinate schemes are guaranteed to produce a monotonically non-increasing objective sequence that converges to a block-coordinatewise stationary point of the original non-convex problem, provided each block's surrogate is a tight upper bound (or, as here, a tight lower bound on a maximized rate) at the current iterate [12].

Chapter 3

Literature Review

This chapter reviews research at the intersection of unmanned aerial vehicle (UAV) communications, energy-aware trajectory design, physical-layer security, and full-duplex (FD) multiuser transmission. The emphasis is on the *joint* design decisions relevant to this thesis: flight trajectory, downlink spatial-division multiple-access (SDMA) beamforming, uplink reception, and artificial-noise (AN) covariance in the presence of a malicious node that can both intercept and jam. Related studies are grouped by their modelling assumptions and optimisation objectives, rather than by publication date, to clarify which mechanisms carry over to the present problem.

3.1 UAV mobility, propulsion energy, and resource allocation

UAV-mounted access points exploit controllable three-dimensional geometry to improve air–ground propagation and to serve users beyond the reach of fixed infrastructure [1,2]. A broad overview of UAV-assisted networks by Wu *et al.* [13] places trajectory control alongside communication, sensing, and network coordination; the review of energy-efficient industrial UAV networks by Zhang *et al.* [14] stresses that flight energy, communication energy, and mission completion requirements should not be evaluated independently. These surveys provide the system-level motivation, but they do not prescribe a secrecy-aware FD transceiver for a moving UAV.

The trajectory is a communications resource because the air–ground channel changes with horizontal position and altitude. Zeng and Zhang [3] developed an energy-efficient trajectory-and- communication design for fixed-wing UAVs, exposing the non-monotonic propulsion power–speed relationship. For the fixed-wing model used in this thesis, $P_{\text{flight}}(v) = c_1 v^3 + c_2/v + c_3$: both excessively high and excessively low speeds are energetically costly.

Consequently, a UAV cannot simply hover near a favourable communication location or fly arbitrarily slowly without accounting for its flight budget. Joint designs for multiple UAVs and multiple users further demonstrate that mobility changes both resource competition and achievable rates [2, 15].

Subsequent studies develop complementary energy-aware UAV designs. Sun *et al.* [16] jointly optimise three-dimensional trajectories and radio resources for *solar-powered* UAVs, introducing energy-arrival and energy-causality considerations absent from a battery-only mission. Wang *et al.* [17] adopt deep reinforcement learning for trajectories in UAV-based Internet-of-Things data collection, illustrating an alternative to explicitly constructed successive-convex-approximation (SCA) surrogates. Cai *et al.* [18] jointly design a three-dimensional UAV path and resources for power-efficient intelligent-reflecting-surface (IRS)-assisted non-orthogonal multiple access (NOMA), while Hashir *et al.* [19] examine a multi-objective performance trade-off for UAV-assisted wireless-powered networks. In each case, trajectory design interacts with a second resource dimension, but neither the specific adversarial FD receiver nor the SDMA secrecy constraints considered here are present.

A related line of work separates flight-path control from other radio-design variables. Zhang *et al.* [20] minimise UAV energy for data collection and wireless power transfer along a *given* path, whereas Xiang *et al.* [21] optimise dynamic antenna steering and beamforming in energy-efficient UAV links. Zhang *et al.* [22] study rotatable-array spatial adaptation in UAV millimetre-wave MIMO communication. These results establish the value of optimising antenna pointing or directivity; however, treating the *path* itself as a coupled variable introduces additional kinematic and propulsion-energy constraints. Related extensions include scheduling UAV data collection under three-dimensional blockage [23], reducing operation time in wireless-powered collection [24], and jointly optimising mobility and resources in an urban integrated sensing and communication (ISAC) network [25]. The energy-minimising secure flight problem in this thesis requires an explicit treatment of both radio power and propulsion power over the complete mission.

3.2 Multiuser beamforming and full-duplex aerial communications

Supporting multiple users in the same time–frequency resources requires careful interference management. Multiple-access techniques proposed for aerial networks range from orthogonal scheduling to NOMA and rate-splitting approaches [4, 10]; SDMA instead uses a multi-antenna array to discriminate users spatially. For a UAV that transmits to several downlink users and receives from several uplink users, beamforming and combining

influence the legitimate SINRs, the interference seen by unintended receivers, and the power consumed by the radio-frequency chains. Early work on large-array SDMA energy efficiency [26] highlights the trade-off between multi-antenna gains and hardware-related power expenditure. That work concerns ground cellular transmission, however, and does not incorporate flight geometry or an adversarial link.

Full-duplex operation allows uplink reception and downlink transmission on one time–frequency resource, but introduces residual self-interference (SI) and additional inter-user interference. Digital cancellation for full-duplex UAVs is investigated in [9]. More broadly, Ng *et al.* [27] formulate power-efficient resource allocation for FD distributed antenna networks; Sun *et al.* [28] study robust power efficiency and secure FD transmission with multiple objectives; and Sun *et al.* [29] examine secure resource allocation in FD multi-antenna multicarrier NOMA systems. These studies provide important precedents for coupling transmit power, information leakage, and residual SI, although the legitimate FD transceiver is not a moving energy-constrained UAV.

FD aerial relaying has also been examined from an AN-aided secrecy-outage perspective [30]; such relay analysis does not jointly optimise a moving UAV's FD multiuser SDMA transceiver and total flight energy.

More recent FD transceiver studies jointly design sensing and communication beams [31], and investigate IRS-assisted SI cancellation [32] or resource allocation with a reflecting surface [33]. In FD covert communication, the receiver's detection uncertainty and information freshness can also influence the preferred transceiver setting [34]. These problems are closely related through their SI coupling, but *covertness* and *secrecy* impose different adversarial performance measures. For the present system, FD is both a legitimate operating mode and an adversarial capability: the UAV must serve simultaneous uplink/downlink transmissions while Eve can jam and intercept in the same slot.

3.3 Trajectory optimisation and physical-layer security for UAVs

The open air–ground channel that helps legitimate communication also allows an eavesdropper to receive transmissions. Physical-layer security (PLS) uses the differences between intended and unintended channels to limit information leakage [5]. In multi-antenna systems, confidential beamforming and AN can reduce the eavesdropper's useful SINR [6]; UAV mobility adds a geometric security lever. Early UAV studies co-design path and power for secure links [35, 36], while robust formulations take uncertain eavesdropper locations or channels into account [37].

Several relevant studies extend these foundations. Cui *et al.* [37] design a robust secure UAV trajectory and transmit-power policy. Wang *et al.* [38] jointly examine placement and cooperative jamming for energy-constrained UAV-assisted secure communications. Li *et al.* [39] address resource allocation in a *multi-UAV, multi-eavesdropper* setting, which adds both coordination and adversarial multiplicity. A companion line of work treats cognitive-radio operating constraints using robust trajectory design [40], and secure trajectory design under uncertain wireless conditions has also been studied directly [41]. In an integrated sensing-and-jamming setting, Wei *et al.* [42] investigate the protection of UAV networks by combining communication with environmental sensing and deliberate interference. These papers show that energy, geometry, and security constraints must be considered jointly, but they do not jointly optimise the FD-SDMA downlink and uplink transceivers against an active FD ground eavesdropper.

Secure UAV links are also studied with additional sensing and mobility functions. Wei *et al.* [43] investigate integrated sensing, navigation, and communication for a UAV network facing a *mobile* eavesdropper; Bastami *et al.* [44] optimise effective secrecy throughput in large-scale uplink rate-splitting-based UAV networks under limited channel feedback. Both examples underscore important departures from a stationary, perfectly-known passive interception channel. In contrast, this thesis isolates the impact of an active, jamming-capable receiver and imposes slot-wise information-leakage constraints, while allowing the legitimate UAV to optimise both transmit and receive processing.

Several works incorporate AN or friendly jamming. Tatar Mamaghani and Hong [45] co-design the UAV path, radio power, and AN allocation; Zhang *et al.* [46] jointly optimise a UAV path and a separate ground jammer's beamforming for robust secrecy, with flight energy treated as a constraint. The UAV-swarm design in [47] explicitly balances propulsion energy and secrecy through a multi-objective formulation. Secure hybrid beamforming in a terrestrial–aerial network with multiple reflecting surfaces [48] introduces additional spatial control, but optimises a different infrastructure and propagation architecture. These findings motivate the use of AN and beamforming as *optimisation variables*, not fixed countermeasures. Their value has to be weighed against the power budget when *energy*, rather than secrecy rate, is the primary objective.

3.4 Energy-aware secure transmission and covert alternatives

It is important to distinguish the objective of minimising mission energy from maximising energy efficiency, which is a ratio of delivered information to consumed energy. Cai *et al.* [49, 50] consider energy-efficient secure UAV communications by co-designing trajec-

tory and radio resources for an OFDMA-based system. These are particularly relevant benchmarks because they include flight energy in a secure mobility-design problem. Their communication model differs from this thesis: the resource allocation relies on orthogonal user separation, the adversarial interception does not include the same FD jamming model, and neither a joint multi-antenna SDMA downlink nor multiuser uplink combining is included. Ng *et al.* [51] also analyse energy-efficient secure OFDMA transmission, but without UAV mobility.

Covert communication is related to, but distinct from, information-theoretic secrecy: it seeks to prevent a detector from reliably establishing that a transmission took place, rather than limiting the amount of confidential information an eavesdropper can decode. UAV studies include three-dimensional placement and power control for covert communication [52], energy-replenishment constraints [53], and covert wireless-powered data collection with latency considerations [54]. Deng *et al.* [55] optimise covert transmission in UAV-assisted NOMA, while [56] jointly designs a UAV trajectory and beamforming in an ISAC setting. Energy-efficient collaborative aerial covert links [57] and robust covert ISAC energy efficiency [58] extend this intersection. These studies are relevant to trajectory–radio coupling and adversarial modelling, but their detection-error constraints cannot be substituted directly for the per-user secrecy thresholds imposed here.

3.5 Passive, active, and full-duplex eavesdropper models

In conventional PLS models Eve is a passive receiver, potentially with uncertain channel or position, but without simultaneous deliberate interference [35,37,59]. For example, Zheng *et al.* [59] jointly optimise aerial MIMO transceiver parameters and hovering position in a passive wiretap setting. This is an informative benchmark for exploiting spatial degrees of freedom, but it does not capture a moving full-duplex UAV with simultaneous uplink/downlink transmissions and an actively jamming Eve.

An *active* eavesdropper can both listen and disrupt the legitimate transmission. Mukherjee and Swindlehurst [7,8] analyse FD interception and jamming in multi-antenna systems. Wu *et al.* [60] investigate secure massive-MIMO transmission in the presence of an active eavesdropper. That study establishes the importance of modelling the adversary's actions explicitly rather than treating Eve as an ordinary passive receiver. Secure beamforming against a *multi-antenna active eavesdropper* remains an active research direction, as illustrated by the joint RIS and beamforming design of Han *et al.* [61]. The present thesis differs in focusing on energy-minimising joint UAV trajectory, FD uplink/downlink transceivers, and slot-wise leakage under a specified self-interfering jamming receiver. A

power-constrained eavesdropper is considered in [62] for probabilistic transmissions; its resource-constrained threat model is related, but it does not include the complete aerial FD trajectory–transceiver configuration.

In a UAV-specific model, Liu *et al.* [63] analyse secrecy and transmission outage with an FD adversary able to jam and intercept simultaneously. This illustrates the effect of residual SI in the *eavesdropper's own receiver*, but its analysis does not perform the full trajectory–multiuser-SDMA transceiver optimisation studied here. Other surveillance research places FD jamming/interception capabilities on an *authorised monitoring* node: [64] optimises a legitimate surveillance UAV's placement and jamming action, while [65] uses a UAV-borne fully-connected RIS and antenna selection to improve wireless information surveillance. These works help establish the architectural relevance of joint listening and interference, although their monitoring objectives do not coincide with protecting legitimate communications from an adversary.

A subtle but consequential feature of an FD Eve is that stronger jamming may also increase its residual SI and thereby weaken its own interception SINR. That feedback need not arise in a passive-Eve model, and it does not imply that jamming is always self-defeating: the outcome depends on the adversary's SI cancellation, jamming power, geometry, and desired link. The model adopted in Chapter 4 retains those quantities explicitly so that the resulting conditions can be analysed rather than assumed.

3.6 Robust beamforming, artificial noise, and optimisation methods

Spatially selective interference management is central to a multiuser PLS formulation. Secure beamforming with uncertain channels has been studied for rate-splitting multiuser MISO transmission [66], and for secure integrated sensing and communication [67, 68]. More elaborate adversarial models include information and sensing eavesdroppers in cell-free ISAC [69], spatially correlated channel uncertainty and multi-antenna eavesdroppers in RIS-assisted links [70], and statistical eavesdropper CSI in active intelligent-surface networks [71]. Robust multiuser beamforming with self-sustainable surfaces is considered in [72], while secure reflecting-surface communication under uncertainty is addressed in [73]. These results confirm the importance of spatial and channel assumptions when evaluating AN, but the present thesis uses a simpler single-antenna Eve with known geometric channel, which permits an interpretable closed-form analysis of its jamming effect.

The relative merits of beamforming and AN depend on both geometry and which secrecy constraints are binding. Confidential beams can suppress leakage toward Eve when suf-

ficient transmit spatial degrees of freedom are available; AN may offer a different means of satisfying secrecy constraints but consumes radiated and amplifier power, and may create interference for legitimate receivers unless spatially suppressed. Previous power-efficient secure beamforming work [28, 29, 66] shows the need to optimise this allocation instead of regarding all available radio power as free. The combination of an FD adversary's residual SI, the UAV's geometry, and a total-energy objective motivates a more specific question: *when is nonzero AN necessary, beneficial, or infeasible under the power constraints?* Chapter 5 examines these regimes for the adopted model; its derived conditions should not be interpreted as universal for arbitrary multi-antenna Eves or channel uncertainty.

The coupled trajectory, beamformer, AN covariance, and receive combiner produce a non-convex optimisation problem. Semidefinite relaxation (SDR) converts lifted beamforming variables into positive semidefinite matrices, sometimes at the price of losing a rank-one guarantee [11]. Le *et al.* [74] provide a rank-one optimisation framework for certain transmit-beamforming problem classes; whether a particular SDR meets those assumptions requires an explicit proof. SCA constructs tractable local bounds for non-convex rate and geometry expressions, while alternating optimisation updates interdependent variable blocks. General block-successive minimisation theory [12] motivates convergence analysis under appropriate surrogate and regularity conditions, but does not establish global optimality for the entire coupled UAV problem. The algorithm in Chapter 5 follows this local-solution paradigm and separately studies the special structure of the AN decision.

3.7 Synthesis and research gap

Table 3.1 summarises representative strands of research. It separates *the presence of a design component* from the stronger claim that the component has been jointly optimised under the exact threat, energy objective, and simultaneous-service constraints of this thesis.

The research gap is therefore *the intersection of several previously studied dimensions*, not the absence of prior work on any one dimension. First, the thesis minimises *total mission energy* (fixed-wing propulsion plus communication-related energy) subject to simultaneous uplink/downlink QoS, rather than maximising secrecy throughput, spectral efficiency, or an information-per-Joule ratio. Second, it jointly controls a time-varying trajectory, multiuser downlink SDMA beamformers, uplink combiners, and an AN covariance, instead of fixing one or more of those dimensions. Third, the active FD Eve both intercepts and jams, with residual self-interference explicitly retained in the secrecy constraints. The

Table 3.1: Representative studies and their relationship to the present formulation.

Representative studies	**Main design emphasis**	**Distinction from this thesis**
[3, 16]	Propulsion-aware or energy-supplied UAV trajectory design	No active FD Eve and joint secure UL/DL SDMA design
[49, 50]	Energy-efficient secure UAV trajectory and OFDMA resources	Energy-efficiency objective, orthogonal access, passive interception
[38, 39]	Energy-constrained jamming, secure multi-UAV resource allocation	No jointly designed moving FD-SDMA transceiver facing FD Eve
[44, 59]	MIMO or rate-splitting UAV secrecy	Different access, mobility, or adversary assumptions
[28, 29, 31]	FD power control, secure resource allocation, or beamforming	No UAV flight-energy term and matched active-Eve trajectory design
[7, 60, 61]	Active eavesdropping and related countermeasures	No complete propulsion-aware FD-UAV UL/DL-SDMA joint design
[63]	FD active Eve and UAV secrecy/outage analysis	No joint multiuser trajectory and transceiver energy minimisation
[56, 57]	Aerial covert optimisation with mobility or energy coupling	Covertness rather than slot-wise secrecy constraints

resulting problem is formulated in Chapter 4 and addressed by alternating optimisation and local convex approximation in Chapter 5.

Finally, the energy formulation makes the *decision to use AN at all* a substantive question: additional AN can improve a leakage constraint, but its power and possible leakage into the intended channels must be accounted for. The analysis in Chapters 5 and 6 identifies model-specific conditions for AN redundancy and for a positive jamming requirement, and investigates whether the latter can be met with finite available power. These results complement, rather than replace, the established literature on robust secrecy beams, secure aerial trajectories, and active-eavesdropper modelling.

Chapter 4

System Model

The considered system is illustrated in Fig. 4.1. A full-duplex (FD) UAV equipped with N_t transmit and N_r receive antennas flies at fixed altitude H from $\mathbf{q}_I$ to $\mathbf{q}_F$ over a mission of duration T, discretized into N slots of length $\omega = T/N$ with horizontal position $\mathbf{q}[n] \in \mathbb{R}^2$ and velocity $\mathbf{v}[n]$. Within one time–frequency resource it serves K downlink (DL) users at known ground positions $\mathbf{t}_k$ and U uplink (UL) users at $\mathbf{t}_u$, in the presence of a ground eavesdropper (Eve) at $\mathbf{t}_E$. Eve is itself FD and *active*: it transmits a probing/jamming signal of power P_E while intercepting both link directions, and consequently suffers residual self-interference $\phi_E P_E$ — a quantity that, as shown in Sec. 7, fundamentally limits its ability to overhear the uplink.

4.1 Channel model

All air–ground links are line-of-sight with free-space path loss. The DL channel to user k, the UL channel from user u, and the Eve channel are

$$\mathbf{h}_k[n] = \sqrt{\beta_0}\, d_k^{-\alpha/2}[n]\, \tilde{\mathbf{h}}_k[n], \quad \mathbf{g}_u[n] = \sqrt{\beta_0}\, d_u^{-\alpha/2}[n]\, \tilde{\mathbf{g}}_u[n], \quad \mathbf{h}_E[n] = \sqrt{\beta_0}\, d_E^{-\alpha/2}[n]\, \tilde{\mathbf{h}}_E[n], \tag{4.1}$$

where $d_x[n] = \sqrt{H^2 + \|\mathbf{q}[n] - \mathbf{t}_x\|^2}$ is the 3-D distance, α the path-loss exponent, and β_0 the reference channel gain. Here $\tilde{\mathbf{h}}_k[n], \tilde{\mathbf{h}}_E[n] \in \mathbb{C}^{N_t}$ are unit-norm *transmit* steering vectors and $\tilde{\mathbf{g}}_u[n] \in \mathbb{C}^{N_r}$ a unit-norm *receive* steering vector of the UAV's uniform linear arrays,

$$[\tilde{\mathbf{h}}]_m = \tfrac{1}{\sqrt{M}}\, e^{-j\pi(m-1)\sin\theta[n]}, \qquad m = 1, \ldots, M, \tag{4.2}$$

with $M = N_t$ (transmit) or $M = N_r$ (receive), parameterized by the azimuth angle $\theta[n]$ of the ground node as seen from the UAV (a planar-array simplification that captures the direction dependence relevant to beamforming nulls). The channel from Eve's transmitter to the UAV receive array, $\mathbf{h}_{E,\mathrm{UL}}[n] \in \mathbb{C}^{N_r}$, follows the same model with the receive

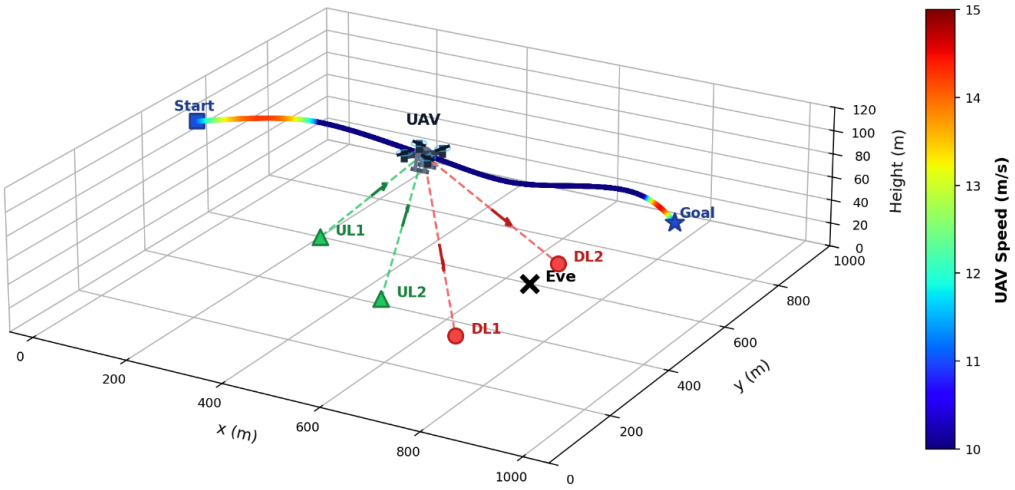


Figure 4.1: System model: a full-duplex UAV with N_t transmit and N_r receive antennas flies from $\mathbf{q}_I$ to $\mathbf{q}_F$ at altitude H, simultaneously serving K downlink and U uplink ground users over LoS channels ($\mathbf{h}_k$, $\mathbf{g}_u$), while a full-duplex active eavesdropper overhears both the downlink ($\mathbf{h}_E$) and the uplink ($g_{E,u}$).

steering vector evaluated at Eve's angle of arrival. Ground links are position-independent: $g_{E,u}$ from UL user u to Eve, and $h_{u\leftrightarrow k}$ from UL user u to DL user k (co-channel cross interference inherent to FD operation). The UAV's transmit array couples into its own receive array through the residual self-interference (RSI) channel $\mathbf{H}_{\mathrm{SI}} \in \mathbb{C}^{N_r \times N_t}$ with suppression factor κ_{SI}.

4.2 Signal model and SINRs

In slot n the UAV transmits $\sum_k \mathbf{w}_k[n]s_k[n]+\mathbf{z}[n]$, where $\mathbf{w}_k$ is the information beamformer for DL user k and $\mathbf{z} \sim \mathcal{CN}(\mathbf{0}, \mathbf{Q}_r[n])$ is artificial noise (AN) with covariance $\mathbf{Q}_r[n] \succeq \mathbf{0}$. The total DL transmit covariance is

$$\mathbf{Q}_{DL}[n] = \textstyle\sum_{k=1}^{K} \mathbf{w}_k[n]\mathbf{w}_k^{\mathsf{H}}[n] + \mathbf{Q}_r[n], \qquad P_{\mathrm{tx}}[n] = \mathrm{Tr}\big(\mathbf{Q}_{DL}[n]\big). \tag{4.3}$$

Legitimate DL. The SINR of DL user k is

$$\gamma_k^{\mathrm{DL}}[n] = \frac{|\mathbf{h}_k^{\mathsf{H}}[n]\mathbf{w}_k[n]|^2}{\sum_{i\neq k} |\mathbf{h}_k^{\mathsf{H}}[n]\mathbf{w}_i[n]|^2 + \mathbf{h}_k^{\mathsf{H}}[n]\mathbf{Q}_r[n]\mathbf{h}_k[n] + \sum_u P_u|h_{u\leftrightarrow k}|^2 + \sigma^2}, \tag{4.4}$$

whose denominator collects intra-cell interference, the UAV's own AN, UL-to-DL cross interference, and noise.

Legitimate UL. With unit-norm receive combiner $\mathbf{f}_u[n]$, the SINR of UL user u is

$$\gamma_u^{\mathrm{UL}}[n] = \frac{P_u\,|\mathbf{f}_u^{\mathsf{H}}[n]\mathbf{g}_u[n]|^2}{\mathbf{f}_u^{\mathsf{H}}[n]\,\mathbf{C}_u[n]\,\mathbf{f}_u[n]}, \qquad \mathbf{C}_u[n] = \sum_{j\neq u} P_j \mathbf{g}_j[n]\mathbf{g}_j^{\mathsf{H}}[n] + \mathbf{R}_{\mathrm{RSI}}[n] + P_E\,\mathbf{h}_{E,\mathrm{UL}}[n]\mathbf{h}_{E,\mathrm{UL}}^{\mathsf{H}}[n] + \sigma^2\mathbf{I}, \tag{4.5}$$

where $\mathbf{R}_{\mathrm{RSI}}[n] = \kappa_{\mathrm{SI}}\mathbf{H}_{\mathrm{SI}}\mathbf{Q}_{DL}[n]\mathbf{H}_{\mathrm{SI}}^{\mathsf{H}}$ is the residual self-interference covariance and $\mathbf{h}_{E,\mathrm{UL}}[n] \in \mathbb{C}^{N_r}$ is the channel from Eve's transmitter to the UAV receive array, so the third term is Eve's *active jamming* received at the UAV — the direct mechanism by which the active eavesdropper degrades the legitimate uplink. The jamming term grows as the UAV approaches Eve, and $\mathbf{R}_{\mathrm{RSI}}$ grows with the UAV's own transmit power, jointly coupling the uplink quality to the trajectory and the downlink design.[1]

Eavesdropper. When decoding DL user k's signal, Eve's SINR is

$$\gamma_{E,k}^{\mathrm{DL}}[n] = \frac{|\mathbf{h}_E^{\mathsf{H}}[n]\mathbf{w}_k[n]|^2}{\sum_{j\neq k}|\mathbf{h}_E^{\mathsf{H}}[n]\mathbf{w}_j[n]|^2 + \mathbf{h}_E^{\mathsf{H}}[n]\mathbf{Q}_r[n]\mathbf{h}_E[n] + \phi_E P_E + \sigma_E^2}, \tag{4.6}$$

and when intercepting UL user u,

$$\gamma_{E,u}^{\mathrm{UL}}[n] = \frac{P_u|g_{E,u}|^2}{\sum_{j\neq u} P_j|g_{E,j}|^2 + \mathbf{h}_E^{\mathsf{H}}[n]\mathbf{Q}_{DL}[n]\mathbf{h}_E[n] + \phi_E P_E + \sigma_E^2}. \tag{4.7}$$

In (4.6)–(4.7) the term $\phi_E P_E$ is Eve's own RSI, and in (4.7) the UAV's *entire* downlink radiation $\mathbf{Q}_{DL}$ acts as friendly jamming that masks the uplink at Eve.

Rates and energy. The achievable rate of any link is $R[n] = \log_2(1+\gamma[n])$, and per-user QoS is imposed on the *time-averaged* rates $\bar{R} = \frac{1}{N}\sum_n R[n]$. The fixed-wing propulsion power and total mission energy are

$$P_{\mathrm{flight}}[n] = c_1\|\mathbf{v}[n]\|^3 + \frac{c_2}{\|\mathbf{v}[n]\|} + c_3, \qquad E_{\mathrm{tot}} = \sum_{n=1}^{N}\omega\Big(P_{\mathrm{flight}}[n] + \zeta P_{\mathrm{tx}}[n] + P_C\Big), \tag{4.8}$$

where ζ is the PA inefficiency and P_C the static circuit power.

[1] In the numerical implementation the jamming term is simplified to an isotropic power floor $P_E\beta_0 d_E^{-\alpha}[n]$, i.e. only its path-loss scaling is retained; this is conservative with respect to combiner-based suppression.

4.3 Problem Formulation

The joint design of the trajectory $\{\mathbf{q}[n], \mathbf{v}[n]\}$, DL beamformers $\{\mathbf{w}_k[n]\}$, AN covariances $\{\mathbf{Q}_r[n]\}$, and UL combiners $\{\mathbf{f}_u[n]\}$ is cast as

$$
\begin{aligned}
(\mathrm{P1}): \quad & \min_{\{\mathbf{q},\mathbf{v}\},\{\mathbf{w}_k\},\{\mathbf{Q}_r\},\{\mathbf{f}_u\}} \quad E_{\mathrm{tot}} \\
& \text{s.t.} \quad \mathrm{C1}: \quad \frac{1}{N}\sum_{n=1}^{N} R_k^{\mathrm{DL}}[n] \;\geq\; R_{\min}^{\mathrm{DL}}, \quad \forall k, && \text{(4.9a)} \\
& \qquad \mathrm{C2}: \quad \frac{1}{N}\sum_{n=1}^{N} R_u^{\mathrm{UL}}[n] \;\geq\; R_{\min}^{\mathrm{UL}}, \quad \forall u, && \text{(4.9b)} \\
& \qquad \mathrm{C3\text{-}DL}: \quad \gamma_{E,k}^{\mathrm{DL}}[n] \;\leq\; \Gamma_{\mathrm{th},E}^{\mathrm{DL}}, \quad \forall k,\, n, && \text{(4.9c)} \\
& \qquad \mathrm{C3\text{-}UL}: \quad \gamma_{E,u}^{\mathrm{UL}}[n] \;\leq\; \Gamma_{\mathrm{th},E}^{\mathrm{UL}}, \quad \forall u,\, n, && \text{(4.9d)} \\
& \qquad \mathrm{C4}: \quad P_{\mathrm{tx}}[n] \;\leq\; P_{\max}^{\mathrm{DL}}, \quad \forall n, && \text{(4.9e)} \\
& \qquad \mathrm{C5}: \quad P_{\mathrm{flight}}[n] + \zeta P_{\mathrm{tx}}[n] + P_C \;\leq\; P_{\mathrm{bdg}}, \quad \forall n, && \text{(4.9f)} \\
& \qquad \mathrm{C6}: \quad V_{\min} \;\leq\; \|\mathbf{v}[n]\| \;\leq\; V_{\max}, \quad \forall n, && \text{(4.9g)} \\
& \qquad \mathrm{C7}: \quad \mathbf{q}[n{+}1] = \mathbf{q}[n] + \omega\,\mathbf{v}[n], \quad \mathbf{q}[1] = \mathbf{q}_I, \quad \mathbf{q}[N] = \mathbf{q}_F, && \text{(4.9h)} \\
& \qquad \mathrm{C8}: \quad \|\mathbf{v}[n{+}1] - \mathbf{v}[n]\| \;\leq\; a_{\max}\,\omega, \quad \forall n, && \text{(4.9i)} \\
& \qquad \mathrm{C9}: \quad \mathbf{Q}_r[n] \;\succeq\; \mathbf{0}, \quad \forall n, && \text{(4.9j)} \\
& \qquad \mathrm{C10}: \quad \mathrm{Tr}\big(\mathbf{F}^{\mathsf{H}}[n]\mathbf{H}_{\mathrm{SI}}\mathbf{Q}_{DL}[n]\mathbf{H}_{\mathrm{SI}}^{\mathsf{H}}\mathbf{F}[n]\big) \;\leq\; \Gamma_{\mathrm{SI}}, \quad \forall n, && \text{(4.9k)}
\end{aligned}
$$

where $\mathbf{F}[n] = [\mathbf{f}_1[n], \ldots, \mathbf{f}_U[n]]$ and the secrecy constraints C3 are enforced *per slot*: they cap Eve's instantaneous SINR in every slot, which yields a per-slot leakage guarantee — strictly stronger than an average-secrecy formulation and deliberately asymmetric to the average-rate QoS in C1–C2 (utility may be averaged; security is guaranteed in the worst case).

Problem (P1) is non-convex on several counts: the fractional SINRs inside the log-rates (C1–C2), the quadratic beamforming terms and trajectory-dependent channels in C3, the reciprocal speed term in P_{flight} (C5 and the objective), the non-convex minimum-speed set in C6, and the bilinear coupling of $\mathbf{F}$ and $\mathbf{Q}_{DL}$ in C10. Moreover the implicit rank-one structure $\mathbf{W}_k = \mathbf{w}_k\mathbf{w}_k^{\mathsf{H}}$ is itself non-convex; we adopt semidefinite relaxation (SDR) and lift $\mathbf{W}_k[n] \succeq \mathbf{0}$, dropping the rank constraint. The next section develops a four-block alternating optimization that addresses each source of non-convexity in turn.

Chapter 5

Proposed Algorithm

Problem (P1) couples four variable groups through the SINRs, the RSI constraint, and the trajectory-dependent channels. We solve it by alternating optimization (AO) over four blocks — DL information beamforming (1a), AN covariance (1b), UL combining (2), and trajectory (3) — where each block solves a convex restriction of (P1) with the remaining variables fixed. Figure 5.1 summarizes the framework. This section derives each block; the convergence properties are established in Sec. 6.

5.1 Rate DC reformulation and the common SCA surrogate

Every achievable rate in (P1) has the form $R[n] = \log_2(1 + \gamma[n])$ with $\gamma[n] = S[n]/I[n]$, where $S[n]$ and $I[n]$ denote the useful-signal and interference-plus-noise powers. The rate admits the difference-of-concave (DC) form

$$R[n] = \log_2\big(S[n] + I[n]\big) - \log_2\big(I[n]\big). \tag{5.1}$$

Although each term is concave in (S, I), their difference is in general neither convex nor concave in the design variables; the DC form is used to construct a concave lower bound. Linearizing the subtracted term at the previous iterate $I^{(t)}[n]$,

$$\log_2\big(I[n]\big) \;\le\; \log_2\big(I^{(t)}[n]\big) + \frac{I[n] - I^{(t)}[n]}{I^{(t)}[n]\ln 2}, \tag{5.2}$$

yields the concave surrogate

$$R^{(t)}[n] \;=\; \log_2\big(S[n] + I[n]\big) - \log_2\big(I^{(t)}[n]\big) - \frac{I[n] - I^{(t)}[n]}{I^{(t)}[n]\ln 2} \;\le\; R[n], \tag{5.3}$$

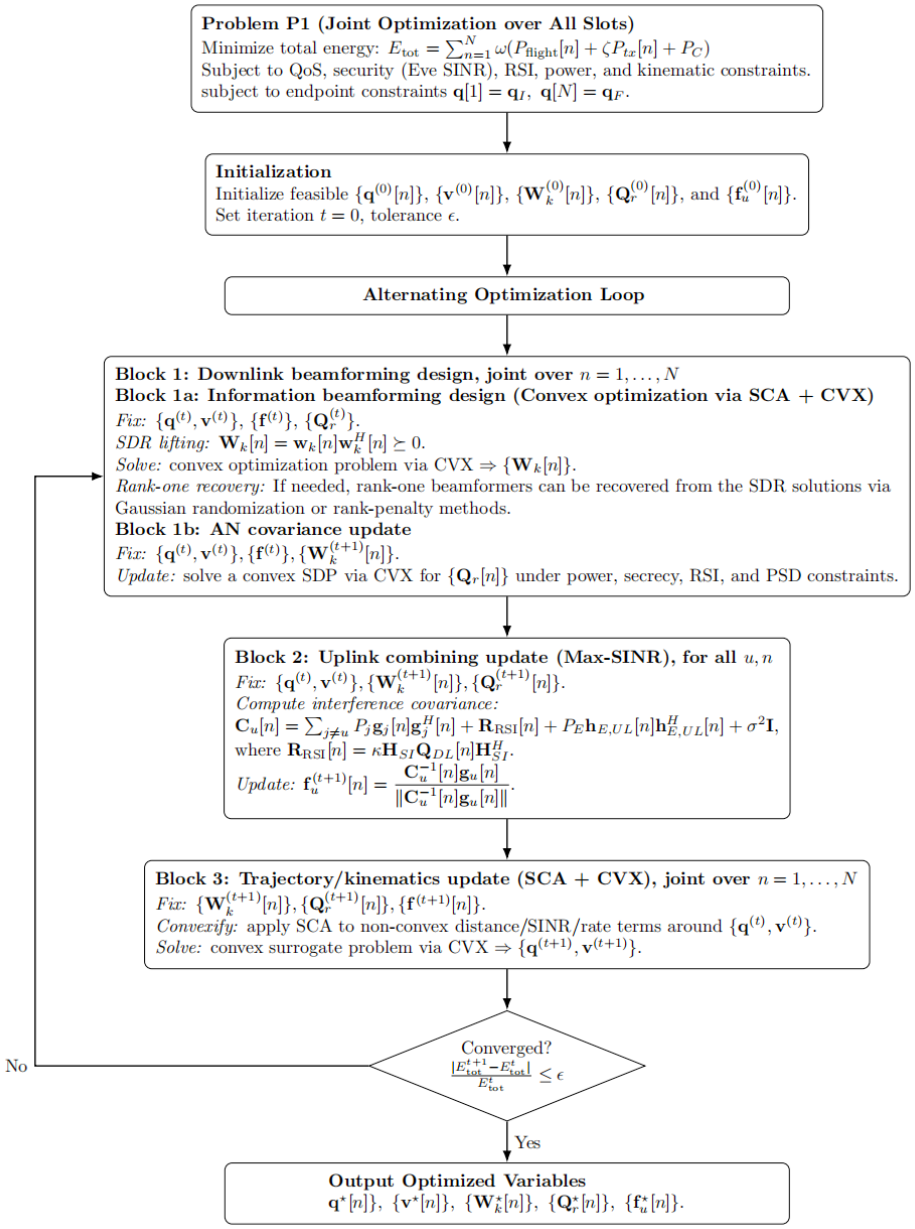


Figure 5.1: AO framework with SDR-based transmit design, closed-form uplink combining, and SCA-based trajectory optimization. The objective is non-increasing over iterations and the loop terminates when the relative improvement falls below the tolerance ϵ.

with equality at $I[n] = I^{(t)}[n]$: the surrogate is a *global under-estimator, tight at the linearization point*. The average QoS constraints C1–C2 are then conservatively enforced as

$$\frac{1}{N}\sum_{n=1}^{N} R^{(t)}[n] \;\geq\; R_{\min}, \tag{5.4}$$

which is the superlevel set of a concave function and hence convex (though not affine), *provided that $S[n]$ and $I[n]$ are affine in the block's optimization variables*. This holds directly in Blocks 1a and 1b, where S and I are linear (trace) functions of the transmit covariances. In Block 3 the same construction applies after one additional linearization step: there S and I depend on the trajectory through the reciprocal path-loss factors, which are first bounded by the affine surrogates of Sec. 5.5 before (5.3) is invoked.

5.2 Block 1a: information beamforming

Fix $\{\mathbf{q}, \mathbf{v}\}$, $\{\mathbf{f}_u\}$, and $\{\mathbf{Q}_r\}$; optimize $\{\mathbf{w}_k[n]\}$. Applying semidefinite relaxation (SDR), lift

$$\mathbf{W}_k[n] = \mathbf{w}_k[n]\mathbf{w}_k^{\mathsf{H}}[n], \qquad \mathbf{W}_k[n] \succeq \mathbf{0}, \;\; \forall k, n, \tag{5.5}$$

and drop the (non-convex) rank-one constraint. Defining $\mathbf{H}_E[n] \triangleq \mathbf{h}_E[n]\mathbf{h}_E^{\mathsf{H}}[n]$, all Eve-side quadratic forms become linear traces:

$$|\mathbf{h}_E^{\mathsf{H}}[n]\mathbf{w}_k[n]|^2 = \mathrm{Tr}\big(\mathbf{H}_E[n]\mathbf{W}_k[n]\big), \qquad \mathbf{h}_E^{\mathsf{H}}[n]\mathbf{Q}_r[n]\mathbf{h}_E[n] = \mathrm{Tr}\big(\mathbf{H}_E[n]\mathbf{Q}_r[n]\big). \tag{5.6}$$

With $\mathbf{Q}_r$ fixed, C3-DL is therefore affine in $\{\mathbf{W}_k\}$:

$$\mathrm{Tr}\big(\mathbf{H}_E[n]\mathbf{W}_k[n]\big) \;\leq\; \Gamma_{\mathrm{th},E}^{\mathrm{DL}}\Big(\sum_{j\neq k}\mathrm{Tr}\big(\mathbf{H}_E[n]\mathbf{W}_j[n]\big) + \mathrm{Tr}\big(\mathbf{H}_E[n]\mathbf{Q}_r[n]\big) + \phi_E P_E + \sigma_E^2\Big), \tag{5.7}$$

and C3-UL, C4, C5, and C10 are likewise affine — C3-UL in the equivalent lower-bound form $\mathrm{Tr}(\mathbf{H}_E[n]\mathbf{Q}_{DL}[n]) \geq c_u$, imposed only when $c_u > 0$ (cf. Sec. 7); C10 via $\mathbf{A}[n] \triangleq \mathbf{H}_{\mathrm{SI}}^{\mathsf{H}}\mathbf{F}[n]\mathbf{F}^{\mathsf{H}}[n]\mathbf{H}_{\mathrm{SI}} \succeq \mathbf{0}$ and the cyclic trace property, $\mathrm{Tr}(\mathbf{A}[n]\mathbf{Q}_{DL}[n]) \leq \Gamma_{\mathrm{SI}}$. The QoS constraints C1–C2 use the surrogate (5.4); note that C2 depends on $\{\mathbf{W}_k\}$ through the residual self-interference $\mathbf{R}_{\mathrm{RSI}}[n]$ in the UL interference (4.5), so the UL rates are also protected during the beamforming update. The resulting subproblem — minimize $\zeta\omega\sum_n \mathrm{Tr}\big(\sum_k \mathbf{W}_k[n]\big)$ (the fixed $\mathrm{Tr}(\mathbf{Q}_r)$ term is constant and omitted) subject to (5.4), (5.7), C3-UL, C4, C5, C9–C10, and $\mathbf{W}_k[n] \succeq \mathbf{0}$ — is convex with PSD matrix variables and is solved by CVX.

Remark (rank-one recovery). Because the SCA-processed QoS constraints are convex but not affine, Block 1a does not belong to the affine-SDP class for which exact rank-one

optimality of SDR is guaranteed. No rank-one optimality claim is made; if rank-one beamformers are required, they can be recovered by Gaussian randomization or rank-penalty refinement, followed by a feasibility restoration step (power rescaling), noting that rescaling interacts with the C3-UL lower bound whenever the latter is active.

5.3 Block 1b: AN covariance update

Fix $\{\mathbf{q}, \mathbf{v}\}$, $\{\mathbf{f}_u\}$, and the newly obtained $\{\mathbf{W}_k\}$; optimize $\{\mathbf{Q}_r[n]\}$. The secrecy constraints (5.7) and the conditional C3-UL lower bound, the power constraints C4–C5, and the RSI constraint C10 are all affine in $\mathbf{Q}_r$ through the lifted forms (5.6). The QoS constraints, however, are *not* affine in $\mathbf{Q}_r$: the AN enters the legitimate users' interference, $\mathbf{h}_k^{\mathsf{H}}\mathbf{Q}_r\mathbf{h}_k$ in (4.4) and $\mathbf{R}_{\mathrm{RSI}}$ in (4.5), inside the log-rates, so C1–C2 again require the DC surrogate (5.4) linearized at the current iterate. Block 1b is therefore a *convex program* (not an affine SDP) with variables $\{\mathbf{Q}_r[n] \succeq \mathbf{0}\}$, solved by CVX.

5.4 Block 2: uplink combiner update

Fix the transmit covariances and trajectory; update the combiners in closed form. The interference-plus-noise covariance seen by UL user u is $\mathbf{C}_u[n]$ from (4.5), and the regularized MMSE update is

$$\mathbf{f}_u[n] \;=\; \frac{\big(\mathbf{C}_u[n] + \lambda_u \mathbf{I}\big)^{-1}\mathbf{g}_u[n]}{\big\|\big(\mathbf{C}_u[n] + \lambda_u \mathbf{I}\big)^{-1}\mathbf{g}_u[n]\big\|}. \tag{5.8}$$

The pure MMSE combiner ($\lambda_u = 0$) maximizes the UL SINR, and the resulting average rate is monotonically non-increasing in λ_u. If the MMSE rate already lies at or below $R_{\min}^{\mathrm{UL}}$ plus a small margin, $\lambda_u = 0$ is kept; otherwise $\lambda_u > 0$ is found by bisection to bring the average UL rate *down* to that level. This deliberate back-off from the MMSE optimum serves numerical stability: driving the UL interference to near zero would push the linearization points $I^{(t)}$ of (5.3) toward zero and destabilize the DC surrogates used in Blocks 1a/1b.

Two conventions adopted in (5.8) matter for the analysis. First, $\|\mathbf{f}_u[n]\| = 1$ is enforced by construction; since the UL SINR is scale-invariant in $\mathbf{f}_u$ while the C10 leakage $\mathrm{Tr}(\mathbf{F}^{\mathsf{H}}\mathbf{H}_{\mathrm{SI}}\mathbf{Q}_{DL}\mathbf{H}_{\mathrm{SI}}^{\mathsf{H}}\mathbf{F})$ scales with $\|\mathbf{f}_u\|^2$, C10 is non-degenerate only under a fixed normalization — without it the constraint could be satisfied vacuously by down-scaling the combiners. Second, (5.8) does not itself enforce C10: after the update, the leakage may transiently violate C10, though the unit-norm constraint bounds the size of this perturbation. In the AO schedule, C10 is enforced exactly in the subsequent Block 1 solve (which treats $\mathbf{F}$ as fixed data), so feasibility with respect to C10 is re-established once per cycle;

see Sec. 6.

5.5 Block 3: trajectory and kinematics

Fix all transceiver variables; optimize $\{\mathbf{q}[n], \mathbf{v}[n]\}$. The trajectory enters (P1) through three channels of influence: the path-loss factors of all air–ground links, the propulsion power, and the kinematic constraints.

Direction-freezing assumption. Within one Block-3 iteration, the steering-vector directions $\tilde{\mathbf{h}}_k, \tilde{\mathbf{g}}_u, \tilde{\mathbf{h}}_E$ are held at their values from the previous iterate and only the path-loss magnitudes vary with $\mathbf{q}[n]$:

$$\beta_x[n] \;\propto\; \big(H^2 + \|\mathbf{q}[n] - \mathbf{t}_x\|^2\big)^{-\alpha/2}. \tag{5.9}$$

This is justified by the trust region below: with $\delta = 80\,\mathrm{m}$ per iteration at altitude $H = 100\,\mathrm{m}$, the per-iteration angular change of any ground node as seen from the UAV is small, so the beamforming inner products $|\tilde{\mathbf{h}}^{\mathsf{H}}\mathbf{w}|$ are approximately constant within the iteration and are refreshed at the next AO cycle.

QoS terms. Introducing slack variables $z_x[n] \geq H^2 + \|\mathbf{q}[n] - \mathbf{t}_x\|^2$ for each link, the reciprocal path-loss terms in the SINRs are handled by the first-order SCA bound

$$\frac{1}{z_x[n]} \;\geq\; \frac{1}{z_x^{(t)}[n]} - \frac{z_x[n] - z_x^{(t)}[n]}{\big(z_x^{(t)}[n]\big)^2}, \tag{5.10}$$

which, combined with the rate surrogate (5.3), yields concave lower bounds on all legitimate rates that are jointly convex in $\{\mathbf{q}, z\}$.

C3-DL: outside-ball (standoff) constraint. With fixed transceivers, define constants $a_k[n], b_k[n], c_E$ such that Eve's DL SINR reads $\gamma_E^{\mathrm{DL}}[n] = a_k[n]|h_E[n]|^2/(b_k[n]|h_E[n]|^2 + c_E)$, where $|h_E[n]|^2 = \beta_0\,(H^2 + \|\mathbf{q}[n] - \mathbf{t}_E\|^2)^{-\alpha/2}$. Imposing $\gamma_E^{\mathrm{DL}}[n] \leq \Gamma_{\mathrm{th},E}^{\mathrm{DL}}$ gives

$$\big(a_k[n] - \Gamma_{\mathrm{th},E}^{\mathrm{DL}} b_k[n]\big)|h_E[n]|^2 \;\leq\; \Gamma_{\mathrm{th},E}^{\mathrm{DL}}\, c_E, \tag{5.11}$$

and two cases arise. *Case 1*: if $a_k[n] \leq \Gamma_{\mathrm{th},E}^{\mathrm{DL}} b_k[n]$, the left-hand side is non-positive and the constraint holds for any position. *Case 2*: otherwise (5.11) is equivalent to the

outside-ball condition

$$H^2 + \|\mathbf{q}[n] - \mathbf{t}_E\|^2 \;\geq\; r_{k,n}^2, \qquad r_{k,n}^2 = \left(\frac{\beta_0 \big(a_k[n] - \Gamma_{\mathrm{th},E}^{\mathrm{DL}} b_k[n]\big)}{\Gamma_{\mathrm{th},E}^{\mathrm{DL}}\, c_E} \right)^{2/\alpha}, \tag{5.12}$$

a non-convex (complement-of-ball) region. Since $d(\mathbf{q}) = \|\mathbf{q} - \mathbf{t}_E\|^2$ is convex, its first-order expansion at $\mathbf{q}^{(t)}[n]$ is a global affine lower bound, and the standard SCA inner approximation

$$H^2 + d\big(\mathbf{q}^{(t)}[n]\big) + 2\big(\mathbf{q}^{(t)}[n] - \mathbf{t}_E\big)^{\mathsf{T}}\big(\mathbf{q}[n] - \mathbf{q}^{(t)}[n]\big) \;\geq\; r_{k,n}^2 \tag{5.13}$$

is affine in $\mathbf{q}[n]$ and tight at the linearization point. *Remark (case switching).* The coefficients $a_k[n], b_k[n]$ are recomputed from the current transceiver iterate at every AO cycle, so a given slot may switch between Case 1 and Case 2 across iterations. This is benign for the convergence analysis: in either case the constraint imposed is a convex inner approximation of the true feasible set, tight at the current iterate, which is the only property the monotonicity argument requires.

C3-UL: inside-ball constraint (conditional). With fixed transceivers, write $\mathbf{h}_E^{\mathsf{H}}[n]\mathbf{Q}_{DL}[n]\mathbf{h}_E[n] = \kappa[n]\,(H^2 + \|\mathbf{q}[n] - \mathbf{t}_E\|^2)^{-\alpha/2}$ with $\kappa[n] \triangleq \beta_0\, \tilde{\mathbf{h}}_E^{\mathsf{H}}\mathbf{Q}_{DL}[n]\tilde{\mathbf{h}}_E \geq 0$ constant in Block 3. For $c_u > 0$ (cf. Sec. 7) the equivalent form of C3-UL becomes

$$H^2 + \|\mathbf{q}[n] - \mathbf{t}_E\|^2 \;\leq\; \Big(\frac{\kappa[n]}{c_u}\Big)^{2/\alpha}, \qquad \forall u, n, \tag{5.14}$$

a convex quadratic (inside-ball) constraint requiring the UAV to remain *close enough* to Eve to mask the uplink. Jointly with (5.12) the feasible region is an annulus around Eve; in all configurations studied in this report $c_u \leq 0$ and (5.14) is absent.

Propulsion and speed constraints. The propulsion power $P_{\text{flight}}[n] = c_1\|\mathbf{v}[n]\|^3 + c_2/\|\mathbf{v}[n]\| + c_3$ is non-convex through the reciprocal term. Introduce $s[n] \geq \|\mathbf{v}[n]\|$, $s[n] > 0$; then $\|\mathbf{v}[n]\|^3 \leq s[n]^3$ handles the cubic term, and for the reciprocal term construct the affine lower bound of the speed at the previous iterate,

$$\underline{v}^{(t)}[n] \;\triangleq\; \|\mathbf{v}^{(t)}[n]\| + \frac{\big(\mathbf{v}^{(t)}[n]\big)^{\mathsf{T}}}{\|\mathbf{v}^{(t)}[n]\|}\big(\mathbf{v}[n] - \mathbf{v}^{(t)}[n]\big) \;\leq\; \|\mathbf{v}[n]\|, \tag{5.15}$$

which is valid for $\mathbf{v}^{(t)}[n] \neq \mathbf{0}$ and is guaranteed positive by the minimum-speed constraint below. Since $x \mapsto 1/x$ is decreasing on $x > 0$, $1/\|\mathbf{v}[n]\| \leq 1/\underline{v}^{(t)}[n]$, and the convex

conservative surrogate of C5 (and of the propulsion term in the objective) is

$$c_1 s[n]^3 + \frac{c_2}{\underline{v}^{(t)}[n]} + c_3 + \zeta P_{\mathrm{tx}}[n] + P_C \;\le\; P_{\mathrm{bdg}}, \qquad \forall n. \tag{5.16}$$

The minimum-speed set $\|\mathbf{v}[n]\| \ge V_{\min}$ in C6 is non-convex (complement of a ball) and is enforced by the same first-order lower bound, $\underline{v}^{(t)}[n] \ge V_{\min}$, which is affine and tight at $\mathbf{v}^{(t)}[n]$; the upper bound $\|\mathbf{v}[n]\| \le V_{\max}$ is a second-order cone constraint. C7 (kinematics with endpoints) and C8 (acceleration) are affine and SOC respectively and are imposed directly.

Trust region. To keep the direction-freezing assumption and all linearizations (5.10), (5.13), (5.15) accurate, the update is confined to $\|\mathbf{q}[n] - \mathbf{q}^{(t)}[n]\| \le \delta$ with $\delta = 80\,\mathrm{m}$. The Block-3 subproblem — minimize the sum of (5.16)-type propulsion surrogates over n subject to the QoS surrogates, (5.13), (5.14) (when active), speed/kinematic constraints, and the trust region — is convex and solved by CVX.

Chapter 6

Convergence Analysis

We now establish the convergence properties of the four-block AO scheme of Chapter 5. Let $E^{(t)}$ denote the objective after iteration t. Each convex block (1a, 1b, 3) solves a convex restriction of (P1) in one variable group with the others fixed, subject to the surrogate constraints of Secs. 5.1–5.5: each surrogate is a global under-estimator of the true rate (resp. an inner approximation of the true feasible set), tight at the current iterate (5.3). Two properties hold for the convex blocks:

1. *Feasibility is preserved:* because the surrogate under-estimates the true rate and is tight at the current iterate, any point feasible for the surrogate problem is feasible for the original problem, and the current iterate is always feasible for the next subproblem.

2. *Monotone descent:* the current iterate is feasible for each block subproblem with objective value $E^{(t)}$, so the block minimizer satisfies $E^{(t+1)} \leq E^{(t)}$.

Block 2 requires a separate remark. It is a closed-form heuristic update, not the solution of a constrained subproblem: it does not include the RSI constraint (C10) and may transiently violate it. However, (i) the combiners do not enter the energy objective, so Block 2 never increases $E^{(t)}$ and the monotonicity argument is unaffected; (ii) $\|\mathbf{f}_u\| = 1$ is enforced by construction, which bounds the RSI perturbation (and is also the normalization convention under which C10 is non-degenerate — without it, scaling $\mathbf{f}_u$ down would satisfy C10 vacuously since the UL SINR is scale-invariant in $\mathbf{f}_u$); and (iii) feasibility is re-established by the subsequent Block 1 solve, which treats $\mathbf{f}$ as fixed data and enforces C10 exactly. The precise convergence statement is therefore: *the objective sequence is monotonically non-increasing and convergent, and feasibility is guaranteed at the exit of every Block 1 (verified numerically: the final* max*-violation is* $\sim 10^{-13}$ *in all experiments)*, rather than pointwise feasibility after every individual block.

Since $E^{(t)}$ is non-increasing and bounded below, the sequence $\{E^{(t)}\}$ converges. Because every surrogate is tight and differentiable at the iterate, any limit point is a *block-coordinatewise stationary point of the SDR-relaxed problem* for the convex blocks, under standard surrogate regularity conditions [12]; no stationarity claim is attached to the heuristic Block 2 beyond the non-increase of the objective. Global optimality cannot be claimed: the AO+SCA scheme is a local method, which is standard for this problem class. Throughout this thesis the trajectory is *initialized as the straight line* from $\mathbf{q}_I$ to $\mathbf{q}_F$ at constant speed, and the trust region ($\delta = 80\,\mathrm{m}$ per iteration) restricts each update to a neighborhood of the current path. All reported solutions are therefore local optima within the attraction basin of the straight-line initialization; absolute energy values should be read with this qualification, whereas the *relative* comparisons between schemes (Section 8.3), which share the same initialization, are unaffected. In our experiments the objective settles within 5–8 iterations (Fig. 8.4).

Chapter 7

Secrecy Analysis

Across every configuration studied in this work, the optimized AN covariance converges to numerical zero, $\mathbf{Q}_r^\star \approx \mathbf{0}$. This chapter shows that this is not an artifact but a structural property of the secrecy constraints, derives an offline-checkable condition for exactly when it holds, and characterizes the complementary regime in which AN becomes mandatory.

7.1 The friendly-jamming threshold c_u

Consider the UL secrecy constraint C3-UL. Substituting Eve's UL interception SINR (4.7) into $\gamma_{E,u}^{\mathrm{UL}}[n] \le \Gamma_{\mathrm{th},E}^{\mathrm{UL}}$ and rearranging,

$$\begin{aligned} P_u|g_{E,u}|^2 \;&\le\; \Gamma_{\mathrm{th},E}^{\mathrm{UL}}\Big(\sum_{j\neq u} P_j|g_{E,j}|^2 + \mathbf{h}_E^{\mathsf{H}}[n]\mathbf{Q}_{DL}[n]\mathbf{h}_E[n] + \phi_E P_E + \sigma_E^2\Big) \\ \iff \quad \mathbf{h}_E^{\mathsf{H}}[n]\mathbf{Q}_{DL}[n]\mathbf{h}_E[n] \;&\ge\; c_u \;\triangleq\; \frac{P_u|g_{E,u}|^2}{\Gamma_{\mathrm{th},E}^{\mathrm{UL}}} - \underbrace{\Big(\sum_{j\neq u} P_j|g_{E,j}|^2 + \phi_E P_E + \sigma_E^2\Big)}_{\beta_u}, \qquad \forall u, n. \end{aligned} \tag{7.1}$$

Equation (7.1) exposes the dual role of the UAV's downlink radiation: the quantity $\mathbf{h}_E^{\mathsf{H}}\mathbf{Q}_{DL}\mathbf{h}_E$ — the total DL power received at Eve — appears as a *lower* bound. When $c_u > 0$, C3-UL is a *friendly-jamming floor*: the UAV must radiate at least c_u toward Eve in *every* slot to mask the uplink. When $c_u \le 0$, the constraint is inherently inactive.

Two properties of c_u are central. First, it is *UAV-independent*: every term in (7.1) is a function of the ground geometry ($g_{E,u}$), the UL transmit powers, Eve's own residual self-interference $\phi_E P_E$, and the secrecy threshold — none of which depend on the trajectory or the beamformers. It can therefore be evaluated *offline*, before any optimization. Second, its sign is dominated by the comparison between the intercepted uplink power $P_u|g_{E,u}|^2$

(scaled by $1/\Gamma^{\mathrm{UL}}_{\mathrm{th},E}$) and Eve's self-interference $\phi_E P_E$.

7.2 The AN-redundancy condition

The friendly-jamming floor (7.1) directly opposes a beamforming null at Eve. If the DL beamformers null Eve's direction, $\mathrm{Tr}(\mathbf{H}_E[n]\mathbf{W}_k[n]) = 0$ for all k, then $\mathbf{h}_E^{\mathsf{H}}\mathbf{Q}_{DL}\mathbf{h}_E = \mathrm{Tr}(\mathbf{H}_E\mathbf{Q}_r)$ and (7.1) would force $\mathbf{Q}_r \neq \mathbf{0}$ whenever $c_u > 0$: the null (no power toward Eve) and the jamming floor (power toward Eve) cannot coexist for free. This tension resolves into the following characterization.

> **AN-redundancy condition.** *The optimal AN vanishes, $\mathbf{Q}_r^\star = \mathbf{0}$, if and only if*
>
> *(i) $c_u \leq 0$ for all u (the UL secrecy constraint is inherently inactive), and*
>
> *(ii) $N_t > K$ with Eve's direction $\tilde{\mathbf{h}}_E$ separable from the users' directions $\{\tilde{\mathbf{h}}_k\}$ (a zero-forcing null at Eve is available without sacrificing the users' beamforming gains).*

Sufficiency. Under (i) the lower bound (7.1) is void, so no power toward Eve is required. Under (ii) the beamformers can satisfy C3-DL by placing $\mathbf{h}_E$ in the null space of every $\mathbf{W}_k$ at zero power cost, so Eve's DL SINR (4.6) vanishes without AN. Any $\mathbf{Q}_r \succ \mathbf{0}$ then strictly increases the objective (through $\zeta\mathrm{Tr}(\mathbf{Q}_r)$) and strictly degrades the legitimate rates (through $\mathbf{h}_k^{\mathsf{H}}\mathbf{Q}_r\mathbf{h}_k$ and $\mathbf{R}_{\mathrm{RSI}}$), so the minimizer sets $\mathbf{Q}_r^\star = \mathbf{0}$.

Necessity. If (i) fails, (7.1) forces $\mathbf{h}_E^{\mathsf{H}}\mathbf{Q}_{DL}\mathbf{h}_E \geq c_u > 0$; a beamforming null then transfers the entire floor onto $\mathrm{Tr}(\mathbf{H}_E\mathbf{Q}_r)$, and relaxing the null merely trades AN power for information-signal leakage toward Eve, which C3-DL caps — either way nonzero power must be radiated at Eve, and if the information beams are to remain protected, part of it must be AN. If (ii) fails, no free null exists and C3-DL must be met by some combination of beam rotation (rate loss), power reduction, or AN in the denominator of (4.6); AN is then typically required whenever the users' QoS constraints pin the beam gains, though the precise trade-off depends on the geometry.

7.3 Verification in the studied configurations

Since c_u is offline-computable, condition (i) is verified directly for every configuration in this thesis (all four narrative scenarios and the benchmark, including the placements with Eve relocated to the UAV corridor at $(650, 600)$). With the baseline parameters ($\phi_E = 0.3$,

$P_E = 0.1\,\mathrm{W}$, $\Gamma^{\mathrm{UL}}_{\mathrm{th},E} = 0.5$),

$$c_u \;\approx\; \underbrace{\frac{P_u|g_{E,u}|^2}{\Gamma^{\mathrm{UL}}_{\mathrm{th},E}}}_{\sim 10^{-9}\,\mathrm{W}} \;-\; \underbrace{\phi_E P_E}_{3\times 10^{-2}\,\mathrm{W}} \;-\; \cdots \;\approx\; -3\times 10^{-2} \;<\; 0 \qquad \forall u, \tag{7.2}$$

i.e. Eve's own residual self-interference exceeds the uplink signal power it intercepts by roughly *seven orders of magnitude*, so (i) holds with enormous margin in all experiments; condition (ii) holds since $N_t = 4 > K = 2$ and Eve's steering vector is distinct from the users'. The observed $\mathbf{Q}_r^\star = \mathbf{0}$ is therefore the true optimum in these configurations, not an artifact.

Physical interpretation: active eavesdropping is self-defeating on the uplink. Equation (7.2) isolates, in closed form, a property of *active* FD eavesdroppers: the probing power P_E that makes Eve a stronger adversary on the downlink simultaneously generates residual self-interference $\phi_E P_E$ that masks the very uplink it tries to overhear. Within this model, Eve's uplink interception is viable only under near-perfect self-interference cancellation ($\phi_E \lesssim 10^{-9}$ at these link budgets). Conversely, a near-passive Eve (small ϕ_E) or UL users located close to Eve drive c_u positive and make AN *mandatory* as friendly jamming.

7.4 The mandatory regime: reachability and a power scaling law

To complete the characterization, consider the regime $c_u > 0$, obtained e.g. with $\phi_E = 10^{-8}$ and $\Gamma^{\mathrm{UL}}_{\mathrm{th},E} = 10^{-2}$, for which $c_u = 1.4\times 10^{-7}$ for the binding UL user in the benchmark geometry. Throughout this section, with slight abuse of notation, $d_E[n] \triangleq \|\mathbf{q}[n] - \mathbf{t}_E\|$ denotes the *horizontal* UAV–Eve distance. The most power-efficient way to satisfy the per-slot floor (7.1) is to align the AN with Eve's direction ($\mathbf{Q}_r \propto \tilde{\mathbf{h}}_E \tilde{\mathbf{h}}_E^{\mathsf{H}}$), giving the minimum AN power

$$\mathrm{Tr}\big(\mathbf{Q}_r\big)_{\min}[n] \;=\; \frac{c_u}{\|\mathbf{h}_E[n]\|^2} \;=\; \frac{c_u\left(H^2 + d_E^2[n]\right)^{\alpha/2}}{\beta_0} \;\overset{\alpha=2}{=}\; \frac{c_u\left(H^2 + d_E^2[n]\right)}{\beta_0}, \tag{7.3}$$

which grows *quadratically* with $d_E[n]$. Evaluating (7.3) along the mission corridor shows the constraint is unreachable at the baseline budget: with $P^{\mathrm{DL}}_{\max} = 3\,\mathrm{W}$, 49 of the 50 slots require between 3.3 and 60 W of AN alone (only the slot directly above Eve is marginally feasible). Since (7.1) must hold in *every* slot while the fixed endpoints $\mathbf{q}_I, \mathbf{q}_F$ keep the UAV far from Eve for most of the mission, the problem is *structurally* infeasible — by

the geometry of the task, not by numerics.

Because (7.3) is analytic, the budget required for reachability follows directly:

$$P_{\max}^{\mathrm{DL}} \gtrsim \frac{c_u \left(H^2 + d_{E,\max}^2\right)}{\beta_0} + P_{\mathrm{serve}}, \tag{7.4}$$

where $d_{E,\max}$ is the largest horizontal UAV–Eve distance the mission *forces* (here the start point, $d_{E,\max} = 650\,\mathrm{m}$) and P_{serve} is the power reserved for the legitimate downlink. Numerically (7.4) gives $\approx 65\,\mathrm{W}$: raising the budget to 50 W leaves only the initial slots infeasible, and $\sim 65\,\mathrm{W}$ makes the mandatory regime formally reachable. However, at the PA inefficiency $\zeta = 5$ this corresponds to $\sim 325\,\mathrm{W}$ of transmit-chain consumption — exceeding the UAV's entire propulsion power. The conclusion is twofold: under per-slot secrecy semantics, (i) the mandatory-jamming regime is unreachable at practical downlink budgets, with the required power scaling as the *square* of the mission-enforced Eve distance (7.4); and (ii) even where formally reachable, per-slot friendly jamming is energetically uneconomical for an energy-constrained UAV, its transmit cost dominating the flight cost itself. This quantifies a hard limit on the anti-eavesdropping capability of the FD system, set jointly by the power budget, the secrecy semantics, and the mission geometry. (Under a *time-averaged* secrecy metric the picture would differ — the UAV could concentrate jamming in the slots near Eve — but averaging weakens the guarantee to statistical leakage and is not adopted here; cf. the per-slot justification in Chapter 4.)

7.5 Numerical caveat

The Eve-direction quadratic forms $\mathbf{h}_E^{\mathsf{H}}\mathbf{X}\mathbf{h}_E$ are $O(10^{-8})$ in physical units, at the tolerance floor of interior-point solvers; constraints of the form (7.1) are silently dropped by SDPT3 unless the Eve direction is normalized (all Eve-side terms divided by $\|\mathbf{h}_E\|^2$, restoring $O(1)$ scale). In the configurations of this thesis this is immaterial because $c_u < 0$ keeps the constraint inactive, but any future configuration operating with $c_u > 0$ must rescale these constraints before trusting the solver output.

Chapter 8

Results

8.1 Simulation Setup

Table 8.1: Baseline parameters.

Parameter	Value	Parameter	Value
T	$100\,\mathrm{s}$	N	50
H	$100\,\mathrm{m}$	ω	$2\,\mathrm{s}$
$N_t = N_r$	4	$K = U$	2
$V_{\min}, V_{\max}$	$10,\ 30\,\mathrm{m/s}$	$a_{\max}$	$5\,\mathrm{m/s^2}$
$\mathbf{q}_I, \mathbf{q}_F$	$(0, 500), (1000, 500)$	$P_{\mathrm{DL}}^{\max}$	$2\,\mathrm{W}$
β_0	10^{-3}	α	2
c_1, c_2, c_3	$9.26 \times 10^{-4}, 2250, 5$	ζ	5
$R_{\min}^{\mathrm{DL}}$	0.3–1.2	$R_{\min}^{\mathrm{UL}}$	0.5

Users are placed to create an asymmetric topology: two DL users (one lower-left, one upper), two UL users, and Eve. The exact demand and Eve position are varied per experiment as described below.

8.2 How the trajectory is shaped

Figure 8.1 shows how the optimal trajectory is progressively reshaped as we add physical factors. The trajectory is colored by instantaneous UAV speed.

- **(a) High transmit power.** With abundant power ($P_{\mathrm{DL}}^{\max} = 50\,\mathrm{W}$) and relaxed QoS, the rate constraint $\log_2(1 + P\beta_0 d^{-\alpha}/\sigma^2) \geq R_{\min}$ can be met by power alone

at any position: the communication constraints no longer shape the trajectory, and the trajectory problem reduces to a pure propulsion problem. From the straight-line initialization the AO then simply retains the straight constant-speed path shown in panel (a). We note explicitly that this straight path is a *local* solution tied to the initialization, not the global optimum of the propulsion problem: under a fixed mission time $T = 100\,\mathrm{s}$ the UAV must remain airborne for the full duration, and the fixed-wing power curve $P(v) = c_1 v^3 + c_2/v + c_3$ is minimized at $v^\star = (c_2/3c_1)^{1/4} = 30\,\mathrm{m/s}$ ($P(30) \approx 105\,\mathrm{W}$), whereas the 1000 m straight line forces $v \approx 10\,\mathrm{m/s}$ ($P(10) \approx 231\,\mathrm{W}$). A longer path flown near $v^\star$ therefore consumes substantially less propulsion energy; a static check confirms such paths are feasible in this scenario. The panel should thus be read as: *with abundant power the trajectory is shaped by propulsion alone*—which is precisely the regime where communication-aware trajectory optimization has no role, the point this baseline is meant to establish.

- **(b) Normal power, uniform demand.** Under a limited power budget the power route above is capped: with $P \le P^{\max}$ fixed, the only remaining lever in the SINR $P\beta_0 d^{-\alpha}/\sigma^2$ is the distance. Meeting $R_{\min}$ requires $d \le \left(P^{\max}\beta_0/(\sigma^2(2^{R_{\min}}-1))\right)^{1/\alpha}$ for enough of the mission, which acts as a soft "service radius" the trajectory must enter. The path therefore bends toward the users, and the bend is deeper the tighter the power budget—this is precisely the regime in which trajectory optimization has value.

- **(c) One high-demand user.** Raising the upper DL user's requirement to $R_{\min} = 1.2$ pulls the trajectory toward that single user, even away from the denser region. Two pieces of theory explain this. First, rate is logarithmic in SINR, so each additional bit of required rate demands an *exponentially* larger channel gain: serving $R_{\min} = 1.2$ instead of 0.3 multiplies the required SINR by $(2^{1.2}-1)/(2^{0.3}-1) \approx 6$, which under $d^{-\alpha}$ path loss with $\alpha = 2$ translates to flying roughly $\sqrt{6} \approx 2.4$ times closer. A single demanding user therefore exerts a stronger "pull" on the trajectory than several easily-served ones. Second, the QoS constraint is on the *time-averaged* rate $\frac{1}{N}\sum_n R_k[n] \ge R_{\min}$, not the per-slot rate. This permits time sharing: the UAV may serve one user intensely in some slots and the other in different slots, provided each average clears its floor. The observed pattern—dwell near the demanding user, a brief excursion to the low-demand user, and speed dropping to $V_{\min}$ near users while rising toward $V_{\max}$ in transit—is exactly the optimal policy for an average-rate constraint: slow flight maximizes the number of high-gain slots near a user, fast flight minimizes time wasted in low-gain regions. The trajectory is thus allocated by *demand weighting*, not user count.

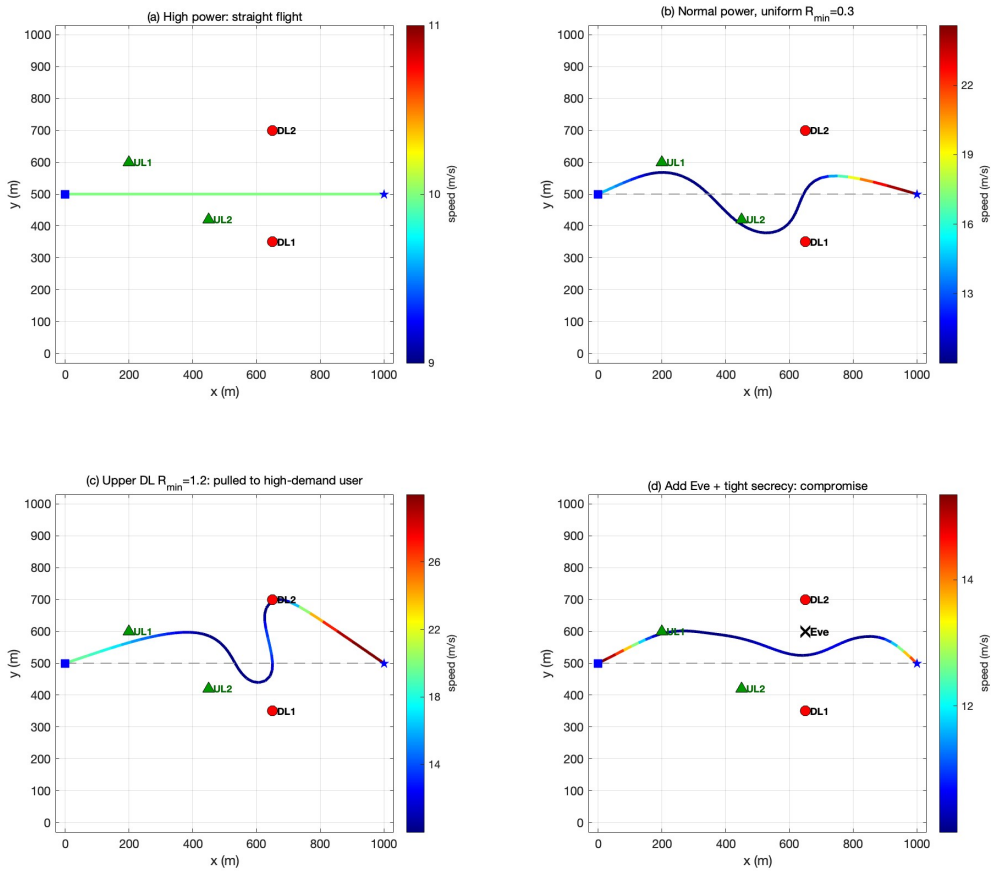


Figure 8.1: Trajectory shaping under increasing physical factors: (a) high power → straight; (b) normal power → bends to users; (c) high-demand user → demand-weighted detour; (d) with Eve → compromise between service and secrecy. Color encodes UAV speed.

- **(d) Adding the eavesdropper.** Placing Eve on the path to the high-demand user and tightening secrecy forces a compromise. The mechanism is that the UAV's position enters *two* distance terms with opposite signs: approaching the user increases the legitimate gain $\|\mathbf{h}_k\|^2 \propto d_k^{-\alpha}$ (good), but because Eve sits on the same approach corridor, it simultaneously increases Eve's gain $\|\mathbf{h}_E\|^2 \propto d_E^{-\alpha}$, tightening the secrecy constraint $\gamma_E \leq \Gamma_E^{\mathrm{DL}}$. Beamforming can null Eve's direction, but the null costs spatial degrees of freedom and transmit power exactly where the user also needs to be served, so the marginal cost of proximity rises sharply inside the corridor. The optimal position balances the two gradients and settles *between* the user and the standoff distance imposed by Eve. The speed profile also *reverses* relative to (b)–(c): the UAV crosses the safe outer region fast and slows once inside the sensitive zone, because there every slot spent near the user must simultaneously respect the secrecy constraint, making dwell time there the scarce resource.

8.3 Benchmark: value of joint optimization

To isolate the contribution of each design dimension, we compare the full joint scheme against two ablations under an identical scenario (uniform $R_{\min}^{\mathrm{DL}} = 0.3$, Eve near the path, $\Gamma_E^{\mathrm{DL}} = 0.2$): *Beamforming Only*, which fixes the trajectory to the initial straight line and optimizes only $\{\mathbf{W}_k, \mathbf{Q}_r, \mathbf{f}_u\}$; and *Trajectory Only*, which fixes the beamformer to the Beamforming-Only solution and optimizes only the path. Results are summarized in Table 8.2.

Table 8.2: Total energy under the safety scenario ($R_{\min}^{\mathrm{DL}} = 0.3$, Eve near path, $\Gamma_E^{\mathrm{DL}} = 0.2$).

Scheme	E_{tot} (kJ)	vs. Beamforming Only
Beamforming Only (fixed straight)	22.8	—
Trajectory Only (fixed beamforming)	21.0	-7.9%
Joint Optimization	**20.7**	$\mathbf{-9.2\%}$

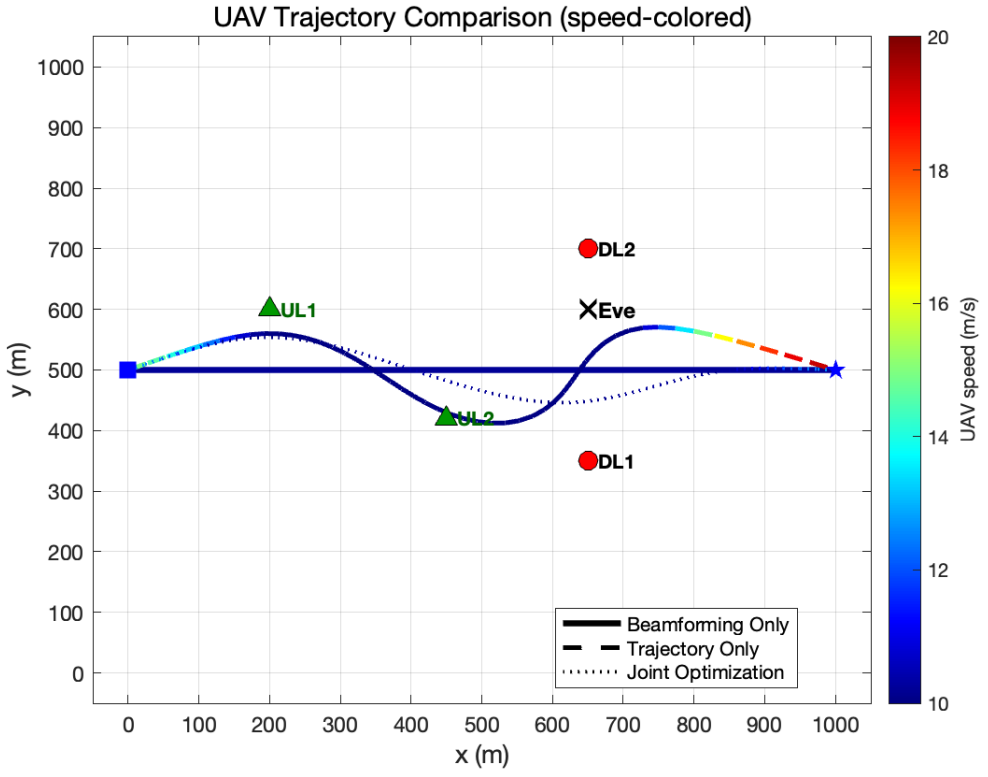


Figure 8.2: Trajectory comparison of the three schemes (speed-colored).

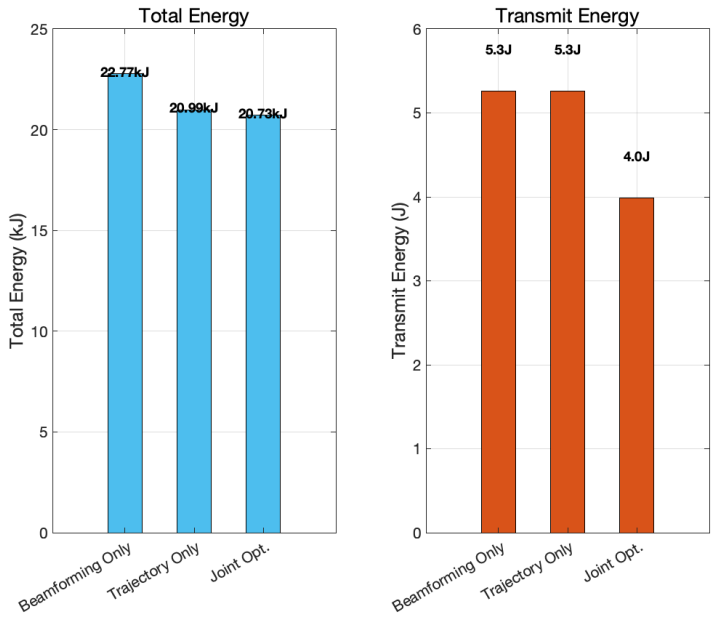


Figure 8.3: Final energy breakdown. Left: total energy (kJ); right: transmit energy (J).

Two observations follow. First, propulsion dominates the energy budget: at the operating speeds, P_{flight} is on the order of 100 W while the amplified transmit power ζP_{tx} is a few watts, so over 99% of E_{tot} is propulsion. In optimization terms, almost the entire gradient of the objective lives in the trajectory variables $(\mathbf{q}, \mathbf{v})$. Beamforming Only removes exactly those variables from the problem, freezing the dominant term at its straight-line value—it can only polish the few-watt transmit term, hence the highest energy. Trajectory Only retains the dominant descent direction and recovers most of the gain (-7.9%). Second, the two design dimensions are *coupled*, not additive: the channels $\mathbf{h}_k(\mathbf{q})$ are functions of the position, so the optimal beamformer depends on the path and vice versa. Trajectory Only inherits a beamformer tuned for the straight line and must fly a path that remains compatible with it, i.e. it optimizes over a restricted feasible set; the joint scheme optimizes over the full set and is therefore guaranteed to do at least as well, which materializes as the additional gain to -9.2%. Relative to its own straight-line starting point, the joint scheme reduces total energy by 12.3% over the AO iterations. The ordering Joint < Trajectory Only < Beamforming Only is thus not accidental but a direct consequence of where the objective's gradient is concentrated and of the channel–trajectory coupling. (All three schemes share the same straight-line initialization, so this comparison is between local solutions in the same attraction basin; the ordering argument above—which variables carry the objective's gradient, and the restriction of the feasible set under a fixed beamformer—does not depend on global optimality.)

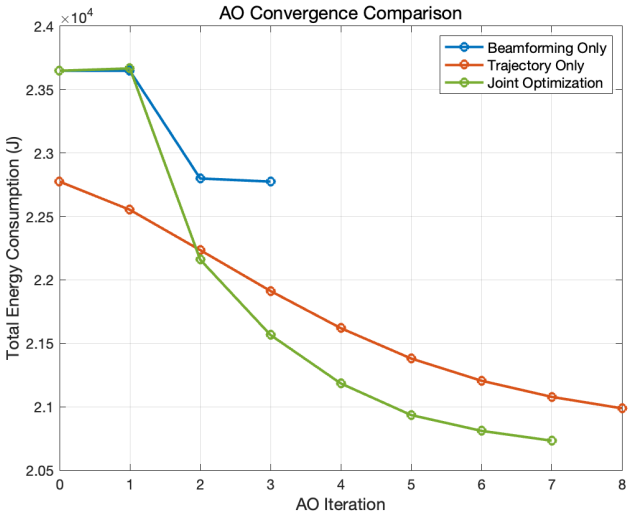


Figure 8.4: AO convergence: total energy decreases monotonically and converges within a few iterations for all three schemes.

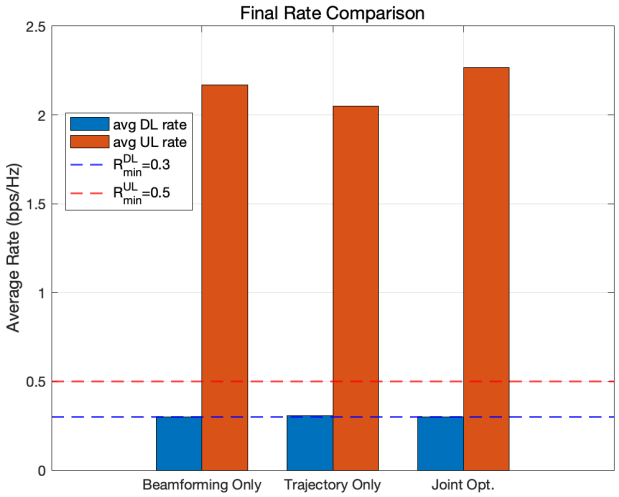


Figure 8.5: Final average DL/UL rates. All schemes meet the QoS floors: the DL rates settle near $R_{\min}^{\mathrm{DL}}$, while the UL rates lie well above $R_{\min}^{\mathrm{UL}}$.

Why the UL rates exceed their floor (Fig. 8.5). This is a textbook active/inactive constraint pattern. At a KKT point, a constraint with a positive Lagrange multiplier is tight, and a constraint satisfied with slack has multiplier zero. The DL rate trades directly against the objective through the transmit term ζP_{tx}: any DL rate above the floor could be converted into energy savings by lowering power, so at the optimum the DL constraints are *active* and settle exactly at $R_{\min}^{\mathrm{DL}}$ (observed [0.315, 0.300]). The UL rate, by contrast, does not appear in the energy objective at all—the UL transmit powers P_u are fixed inputs, not variables—so there is no descent direction that benefits from reducing it; its multiplier is zero and the constraint is *inactive*. The size of the slack is governed jointly

by the physics (the $N_r = 4$ array gain and short UAV–UL distances make a high UL SINR available) and by the Block 2 regularization: the bisected λ_u deliberately holds the average UL rate near $R_{\min}^{\mathrm{UL}} + 0.2$ rather than at the unregularized MMSE maximum, to keep the DC linearization points in Blocks 1a/1b well conditioned. Where the reported UL averages exceed this level, the regularizer was inactive ($\lambda_u = 0$) and the physics alone set the rate. Tightening $R_{\min}^{\mathrm{UL}}$ or lowering P_u would move the UL rates toward the floor but would not change the energy conclusions above.

8.4 Artificial noise in the numerical results

In every configuration reported above — all four narrative scenarios and the benchmark — the optimized AN covariance converges to numerical zero, $\mathbf{Q}_r^\star \approx \mathbf{0}$, typically within the first AO iteration. As established in Chapter 7, the AN-redundancy condition holds in all of these configurations ($c_u \approx -3 \times 10^{-2} < 0$ for every uplink user, and $N_t = 4 > K = 2$ with separable directions), so $\mathbf{Q}_r^\star = \mathbf{0}$ is the true optimum.

Chapter 9

Conclusion

This thesis has developed a joint design of trajectory, SDMA beamforming, artificial-noise allocation, and uplink combining for a full-duplex UAV serving multiple downlink and uplink users under per-user rate and per-slot secrecy constraints, against a ground eavesdropper that is itself full-duplex and active. The design objective — minimizing the UAV's total propulsion and communication energy — departs from the secrecy-rate-maximization framing that dominates the UAV physical-layer-security literature, and this change of objective proved consequential for every conclusion that followed.

The resulting non-convex problem was decomposed into a four-block alternating optimization combining semidefinite relaxation and successive convex approximation. The objective sequence is monotonically non-increasing and convergent, with feasibility guaranteed at the exit of each optimization cycle and limit points that are block-coordinatewise stationary for the relaxed problem under standard surrogate conditions. Numerically, the joint design reduces total energy by 12.3% relative to its straight-line starting point, and by 9.2% relative to a beamforming-only baseline in a representative secrecy-constrained scenario. The ordering Joint $<$ Trajectory Only $<$ Beamforming Only is explained by where the objective's gradient is concentrated: propulsion dominates the energy budget by roughly two orders of magnitude over the amplified transmit power, so the trajectory variables carry almost the entire descent direction, while the channel–trajectory coupling makes the two design dimensions non-additive. The optimized trajectory itself is shaped by an interpretable interplay of forces — transmit-power budgets set a service radius, logarithmic rates make high-demand users pull disproportionately, and the eavesdropper imposes a standoff — each of which was isolated in a controlled numerical narrative.

On the security side, this thesis has derived and verified an offline-checkable condition under which artificial noise is redundant: AN vanishes at the optimum exactly when (i) the friendly-jamming threshold satisfies $c_u \leq 0$ for every uplink user — a property of the ground geometry, the uplink powers, and Eve's own residual self-interference alone,

independent of the trajectory and beamformers — and (ii) the spatial degrees of freedom suffice for a zero-forcing null at Eve ($N_t > K$ with separable directions). The sign of c_u reduces to a comparison between the uplink power Eve intercepts and the self-interference $\phi_E P_E$ generated by its own probing: in the studied configurations the latter exceeds the former by seven orders of magnitude, making active eavesdropping *self-defeating* on the uplink and the observed $\mathbf{Q}_r^\star = \mathbf{0}$ the true optimum rather than an artifact. In the complementary regime ($c_u > 0$), AN becomes a mandatory per-slot friendly-jamming floor whose minimum power grows quadratically with the mission-enforced UAV–eavesdropper distance; this scaling law renders the mandatory regime structurally unreachable at practical downlink budgets, and energetically uneconomical even where formally reachable. Under per-slot secrecy semantics, the anti-eavesdropping capability of the FD system is therefore subject to a hard limit set jointly by the power budget, the secrecy semantics, and the mission geometry.

9.1 Future Work

Multi-antenna eavesdropper. The redundancy condition is tied to a single-antenna Eve under pure-LoS channels, where one spatial null suffices. A multi-antenna Eve observes the transmission through multiple spatial dimensions, so no single beamforming null can suppress its reception; AN then becomes non-trivial on the downlink, and the interplay between the null space dimension $N_t - K$, Eve's antenna count, and the AN covariance structure is an open design problem.

Phase diagram of AN necessity. The AN-redundancy condition is binary in the sign of c_u. Sweeping ϕ_E (and the UL-user-to-Eve distances) across the $c_u = 0$ boundary would turn the conditional result into a characterized phase diagram — mapping, over the parameter space, where AN is redundant, where it is mandatory but unreachable, and where (if anywhere) it is both required and affordable. The numerical rescaling prescribed in Sec. 7.5 is a prerequisite for trustworthy solver output in the $c_u > 0$ region.

Relaxing the eavesdropper-CSI assumption. The design assumes Eve's position, and hence $\mathbf{h}_E[n]$, is known. Sensing-assisted predictive beamforming, in which the aerial link's own echoes are used to track a low-altitude target and to predict the beamformer one slot ahead [75], offers a route to estimating an uncooperative node's direction online. Coupling such a tracking loop to the secrecy constraints C3 — so that the standoff ball around Eve is re-centred as the estimate is refined — would remove the strongest idealization in the present model.

Average-secrecy variants. Under the per-slot secrecy semantics adopted here, the mandatory-jamming regime is unreachable by the geometry of the task. A time-averaged secrecy metric would re-open it — the UAV could concentrate jamming in the slots near Eve — at the cost of weakening the guarantee to statistical leakage. Formalizing this trade-off, and identifying threat models (e.g. long-term traffic analysis rather than instantaneous interception) under which averaging is defensible, would connect the secrecy semantics to the achievable mission set.

Bibliography

[1] Y. Zeng, R. Zhang, and T. J. Lim, "Wireless communications with unmanned aerial vehicles: Opportunities and challenges," *IEEE Communications Magazine*, vol. 54, no. 5, pp. 36–42, 2016.

[2] Y. Zeng, Q. Wu, and R. Zhang, "Accessing from the sky: A tutorial on UAV communications for 5G and beyond," *Proceedings of the IEEE*, vol. 107, no. 12, pp. 2327–2375, 2019.

[3] Y. Zeng and R. Zhang, "Energy-efficient UAV communication with trajectory optimization," *IEEE Transactions on Wireless Communications*, vol. 16, no. 6, pp. 3747–3760, 2017.

[4] W. Jaafar, S. Naser, S. Muhaidat, P. C. Sofotasios, and H. Yanikomeroglu, "Multiple access in aerial networks: From orthogonal and non-orthogonal to rate-splitting," *IEEE Open Journal of Vehicular Technology*, vol. 1, pp. 372–392, 2020.

[5] A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, no. 8, pp. 1355–1387, 1975.

[6] S. Goel and R. Negi, "Guaranteeing secrecy using artificial noise," *IEEE Transactions on Wireless Communications*, vol. 7, no. 6, pp. 2180–2189, 2008.

[7] A. Mukherjee and A. L. Swindlehurst, "A full-duplex active eavesdropper in MIMO wiretap channels: Construction and countermeasures," in *Proc. 45th Asilomar Conference on Signals, Systems and Computers*, 2011, pp. 265–269.

[8] ——, "Jamming games in the MIMO wiretap channel with an active eavesdropper," *IEEE Transactions on Signal Processing*, vol. 61, no. 1, pp. 82–91, 2013.

[9] L. Tian, C. Shi, and Z. Xu, "Digital self-interference cancellation for full-duplex UAV communication system over time-varying channels," *Drones*, vol. 7, no. 3, p. 151, 2023.

[10] Z. Wei, D. Xu, S. Li, S. Song, D. W. K. Ng, and G. Caire, "Resource allocation design for next-generation multiple access: A tutorial overview," *Proceedings of the IEEE*, vol. 112, no. 9, pp. 1230–1263, 2024.

[11] Z.-Q. Luo, W.-K. Ma, A. M.-C. So, Y. Ye, and S. Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 20–34, 2010.

[12] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1126–1153, 2013.

[13] Q. Wu, J. Xu, Y. Zeng, D. W. K. Ng, N. Al-Dhahir, R. Schober, and A. L. Swindlehurst, "A Comprehensive Overview on 5G-and-Beyond Networks with UAVs: From Communications to Sensing and Intelligence," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 10, pp. 2912–2945, 2021.

[14] Y. Zhang, R. Zhao, D. Mishra, and D. W. K. Ng, "A Comprehensive Review of Energy-Efficient Techniques for UAV-Assisted Industrial Wireless Networks," *Energies*, vol. 17, 2024.

[15] Q. Wu, Y. Zeng, and R. Zhang, "Joint trajectory and communication design for multi-UAV enabled wireless networks," *IEEE Transactions on Wireless Communications*, vol. 17, no. 3, pp. 2109–2121, 2018.

[16] Y. Sun, D. Xu, D. W. K. Ng, L. Dai, and R. Schober, "Optimal 3D-Trajectory Design and Resource Allocation for Solar-Powered UAV Communication Systems," *IEEE Transactions on Communications*, vol. 67, no. 6, pp. 4281–4298, 2019.

[17] Y. Wang, Z. Gao, J. Zhang, X. Cao, Y. Gao, D. W. K. Ng, and M. D. Renzo, "Trajectory Design for UAV-Based Internet-of-Things Data Collection: A Deep Reinforcement Learning Approach," *IEEE Internet of Things Journal*, vol. 9, no. 5, pp. 3899–3912, 2022.

[18] Y. Cai, Z. Wei, S. Hu, C. Liu, D. W. K. Ng, and J. Yuan, "Resource Allocation and 3D Trajectory Design for Power-Efficient IRS-Assisted UAV-NOMA Communications," *IEEE Transactions on Wireless Communications*, vol. 21, no. 12, pp. 10 315–10 334, 2022.

[19] S. M. Hashir, A. Mehrabi, M. R. Mili, M. J. Emadi, D. W. K. Ng, and I. Krikidis, "Performance Trade-off in UAV-Aided Wireless-Powered Communication Networks via Multi-objective Optimization," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 12, pp. 13 430–13 435, 2021.

[20] J. Zhang, G. Lu, L. Xiang, X. Ge, and D. W. K. Ng, "Energy minimization for UAV-aided data collection along a fixed flight path with a directional antenna," *IEEE Transactions on Communications*, vol. 74, pp. 764–780, 2026.

[21] L. Xiang, X. Zhu, F. Pei, and D. W. K. Ng, "Energy-efficient dynamic array-steering and beamforming for UAV-aided communications," in *Proc. IEEE Global Communications Conference (GLOBECOM)*, 2024.

[22] X. Zhang, L. Xiang, J. Wang, X. Gao, D. W. K. Ng, and R. Schober, "Rotatable Antenna Array Enabled UAV mmWave Massive MIMO Communication," *IEEE Transactions on Communications*, vol. 74, pp. 1219–1236, 2026.

[23] Y. Cai, W. Yuan, Z. Wei, C. Liu, S. Hu, and D. W. K. Ng, "Trajectory Design and Resource Allocation for UAV-Enabled Data Collection in Wireless Sensor Networks with 3D Blockages," in *1st International Conference on 6G Networking (6GNet 2022)*, 2022.

[24] Y. Zhang, D. Mishra, H. H. Gharakheili, and D. W. K. Ng, "UAV Operation Time Minimization for Wireless-Powered Data Collection," in *2024 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2024)*, 2024.

[25] Z. Cao, J. Sun, J. Zhang, and D. W. K. Ng, "Joint trajectory optimization and resource allocation of UAV in urban ISAC network," in *Proc. 30th Asia-Pacific Conference on Communications (APCC)*, 2025.

[26] D. W. K. Ng, E. S. Lo, and R. Schober, "Energy-Efficient Resource Allocation in SDMA Systems with Large Numbers of Base Station Antennas," in *IEEE International Conference on Communications (ICC) 2012, Ottawa, Canada*, 2012.

[27] D. W. K. Ng, Y. Wu, and R. Schober, "Power Efficient Resource Allocation for Full-Duplex Radio Distributed Antenna Networks," *IEEE Transactions on Wireless Communications*, vol. 15, no. 4, pp. 2896–2911, 2016.

[28] Y. Sun, D. W. K. Ng, J. Zhu, and R. Schober, "Multi-Objective Optimization for Robust Power Efficient and Secure Full-Duplex Wireless Communication Systems," *IEEE Transactions on Wireless Communications*, vol. 15, no. 8, pp. 5511–5526, 2016.

[29] ——, "Robust and Secure Resource Allocation for Full-Duplex MISO Multicarrier NOMA Systems," *IEEE Transactions on Communications*, vol. 66, no. 9, pp. 4119–4137, 2018.

[30] T. T. Nguyen, V. T. Hoang, T. T. H. Le, and X. N. Tran, "Physical layer security for UAV-based full-duplex relay NOMA system," in *Proc. International Conference on Advanced Technologies for Communications (ATC)*, 2022.

[31] Z. He, W. Xu, H. Shen, D. W. K. Ng, Y. Eldar, and X. You, "Full-Duplex Communication for ISAC: Joint Beamforming and Power Optimization," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 9, pp. 2920–2936, 2023.

[32] C. Qiu, Q. Wu, M. Hua, W. Chen, S. Ma, F. Hou, D. W. K. Ng, and A. L. Swindlehurst, "Intelligent Reflecting Surface Empowered Self-Interference Cancellation in Full-Duplex Systems," *IEEE Transactions on Communications*, vol. 72, no. 5, pp. 2945–2958, 2024.

[33] D. Xu, X. Yu, Y. Sun, D. W. K. Ng, and R. Schober, "Resource Allocation for IRS-assisted Full-Duplex Cognitive Radio Systems," *IEEE Transactions on Communications*, vol. 68, no. 12, pp. 7376–7394, 2020.

[34] C. Wang, Z. Li, and D. W. K. Ng, "Covert Rate Optimization of Millimeter Wave Full-Duplex Communications," *IEEE Transactions on Wireless Communications*, vol. 21, no. 5, pp. 2844–2861, 2022.

[35] G. Zhang, Q. Wu, M. Cui, and R. Zhang, "Securing UAV communications via trajectory optimization," in *Proc. IEEE Global Communications Conference (GLOBECOM)*, 2017.

[36] ——, "Securing UAV communications via joint trajectory and power control," *IEEE Transactions on Wireless Communications*, vol. 18, no. 2, pp. 1376–1389, 2019.

[37] M. Cui, G. Zhang, Q. Wu, and D. W. K. Ng, "Robust trajectory and transmit power design for secure UAV communications," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 9, pp. 9042–9046, 2018.

[38] W. Wang, X. Li, M. Zhang, K. Cumanan, D. W. K. Ng, G. Zhang, J. Tang, and O. A. Dobre, "Energy-Constrained UAV-Assisted Secure Communications with Position Optimization and Cooperative Jamming," *IEEE Transactions on Communications*, vol. 68, no. 7, pp. 4476–4489, 2020.

[39] R. Li, Z. Wei, L. Yang, D. W. K. Ng, J. Yuan, and J. An, "Resource Allocation for Secure Multi-UAV Communication Systems with Multi-Eavesdropper," *IEEE Transactions on Communications*, vol. 68, no. 7, pp. 4490–4506, 2020.

[40] Y. Zhou, H. Zhou, F. Zhou, D. W. K. Ng, and R. Q. Hu, "Robust Chance-Constrained Trajectory and Transmit Power Optimization for UAV-Enabled CR Networks," in *IEEE International Conference on Communications (ICC)*, 2020.

[41] X. Sun, C. Shen, D. W. K. Ng, and Z. Zhong, "Robust Trajectory and Resource Allocation Design for Secure UAV-aided Communications," in *IEEE International Conference on Communications (ICC)*, 2019.

[42] Z. Wei, F. Liu, D. W. K. Ng, and R. Schober, "Safeguarding UAV Networks through Integrated Sensing, Jamming, and Communications," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.

[43] Z. Wei, F. Liu, C. Liu, Z. Yang, D. W. K. Ng, and R. Schober, "Integrated Sensing, Navigation, and Communication for Secure UAV Networks with a Mobile Eavesdropper," *IEEE Transactions on Wireless Communications*, vol. 23, no. 7, pp. 7060–7078, 2024.

[44] H. Bastami, H. Behroozi, M. Moradikia, A. Abdelhadi, D. W. K. Ng, and L. Hanzo, "Large-Scale Rate-Splitting Multiple Access in Uplink UAV Networks: Effective Secrecy Throughput Maximization Under Limited Feedback Channel," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 7, pp. 9267–9280, 2023.

[45] M. Tatar Mamaghani and Y. Hong, "Joint trajectory and power allocation design for secure artificial noise aided UAV communications," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 3, pp. 2850–2855, 2021.

[46] Y. Zhang, X. Gao, M. Shi, H. Yuan, J. Kang, D. Niyato, and K. Yang, "Robust secure UAV communications with the aid of jamming beamforming," *IEEE Transactions on Communications*, 2025, early access, doi: 10.1109/TCOMM.2025.3566977.

[47] G. Sun, J. Li, A. Wang, Q. Wu, Z. Sun, and Y. Liu, "Secure and energy-efficient UAV relay communications exploiting collaborative beamforming," *IEEE Transactions on Communications*, vol. 70, no. 8, pp. 5401–5415, 2022.

[48] Y. Sun, K. An, Y. Zhu, G. Zheng, K.-K. Wong, S. Chatzinotas, D. W. K. Ng, and D. Guan, "Energy-Efficient Hybrid Beamforming for Multilayer RIS-Assisted Secure Integrated Terrestrial-Aerial Networks," *IEEE Transactions on Communications*, vol. 70, no. 6, pp. 4189–4210, 2022.

[49] Y. Cai, Z. Wei, R. Li, D. W. K. Ng, and J. Yuan, "Energy-efficient resource allocation for secure UAV communication systems," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, 2019.

[50] ——, "Joint trajectory and resource allocation design for energy-efficient secure UAV communication systems," *IEEE Transactions on Communications*, vol. 68, no. 7, pp. 4536–4553, Jul. 2020.

[51] D. W. K. Ng, E. S. Lo, and R. Schober, "Energy-Efficient Resource Allocation for Secure OFDMA Systems," *IEEE Transactions on Vehicular Technology*, vol. 61, pp. 2572–2585, 2012.

[52] X. Zhou, S. Yan, D. W. K. Ng, and R. Schober, "Three-Dimensional Placement and Transmit Power Design for UAV Covert Communications," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 12, pp. 13 424–13 429, 2021.

[53] Y. Wang, S. Yan, X. Zhou, Y. Huang, and D. W. K. Ng, "Covert Communication with Energy Replenishment Constraints in UAV Networks," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 9, pp. 10 143–10 148, 2022.

[54] X. Lu, W. Yang, S. Yan, Z. Li, and D. W. K. Ng, "Covertness and Timeliness of Data Collection in UAV-Aided Wireless-Powered IoT," *IEEE Internet of Things Journal*, vol. 9, no. 14, pp. 12 573–12 587, 2022.

[55] D. Deng, S. Dang, X. Li, D. W. K. Ng, and A. Nallanathan, "Joint Optimization for Covert Communications in UAV-Assisted NOMA Networks," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 1, pp. 1012–1026, 2024.

[56] D. Deng, W. Zhou, X. Li, D. B. da Costa, D. W. K. Ng, and A. Nallanathan, "Joint beamforming and UAV trajectory optimization for covert communications in ISAC networks," *IEEE Transactions on Wireless Communications*, vol. 24, no. 2, pp. 1016–1030, 2025.

[57] J. Wu, W. Yuan, Q. Tao, H. Huang, and D. W. K. Ng, "Aerial-networked ISAC-empowered Collaborative Energy-Efficient Covert Communications," *Chinese J. Aeronautics*, vol. 38, no. 10, 2025.

[58] D. Deng, X. Li, S. Dang, D. W. K. Ng, A. Nallanathan, and D. Niyato, "Energy Efficiency Optimization for Robust Covert ISAC Systems," *IEEE Journal on Selected Areas in Communications*, vol. 44, pp. 562–577, 2026.

[59] Z. Zheng, X. Wang, Z. Fei, Q. Wu, B. Li, and L. Hanzo, "Secure UAV-to-ground MIMO communications: Joint transceiver and location optimization," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 11, pp. 11 661–11 675, 2022.

[60] Y. Wu, R. Schober, D. W. K. Ng, C. Xiao, and G. Caire, "Secure Massive MIMO Transmission with an Active Eavesdropper," *IEEE Transactions on Information Theory*, vol. 62, no. 7, pp. 3880–3900, 2016.

[61] J. Han, L. Yang, Y. Zhan, W. Wang, and D. W. K. Ng, "Joint Beamforming and RIS Design Against a Multi-Antenna Active Eavesdropper," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2026.

[62] Z. Chen, S. Yan, L. Li, L. Yang, D. W. K. Ng, and R. Schober, "Securing Probabilistic Wireless Transmissions Against a Power-Constrained Eavesdropper," in *IEEE International Conference on Communications (ICC)*, 2025.

[63] C. Liu, J. Lee, and T. Q. S. Quek, "Safeguarding UAV communications against full-duplex active eavesdropper," *IEEE Transactions on Wireless Communications*, vol. 18, no. 6, pp. 2919–2931, 2019.

[64] G. Hu, Z. Li, J. Si, K. Xu, D. Xu, Y. Cai, and N. Al-Dhahir, "Maxmin fairness for UAV-enabled proactive eavesdropping with jamming over distributed transmit beamforming-based suspicious communications," *IEEE Transactions on Communications*, vol. 71, no. 3, pp. 1595–1613, 2023.

[65] S. Lin, Y. Zou, H. Li, B. Li, and D. W. K. Ng, "UAV-borne FC-RIS empowered wireless information surveillance with threshold-based antenna selection," *IEEE Transactions on Wireless Communications*, vol. 25, pp. 6160–6176, 2026.

[66] H. Fu, S. Feng, W. Tang, and D. W. K. Ng, "Robust Secure Beamforming Design for Two-User Downlink MISO Rate-Splitting Systems," *IEEE Transactions on Wireless Communications*, vol. 19, no. 12, pp. 8351–8365, 2020.

[67] Z. Ren, L. Qiu, J. Xu, and D. W. K. Ng, "Robust Transmit Beamforming for Secure Integrated Sensing and Communication," *IEEE Transactions on Communications*, vol. 71, no. 9, pp. 5549–5564, 2023.

[68] D. Xu, X. Yu, D. W. K. Ng, A. Schmeink, and R. Schober, "Robust and secure resource allocation for ISAC systems: A novel optimization framework for variable-length snapshots," *IEEE Transactions on Communications*, vol. 70, no. 12, pp. 8196–8214, 2022.

[69] Z. Ren, J. Xu, L. Qiu, and D. W. K. Ng, "Secure Cell-Free Integrated Sensing and Communication in the Presence of Information and Sensing Eavesdroppers," *IEEE Journal on Selected Areas in Communications*, vol. 42, pp. 3217–3231, 2024.

[70] Z. Li, Q. Lin, Y.-C. Wu, D. W. K. Ng, and A. Nallanathan, "Enhancing Physical Layer Security with RIS under Multi-Antenna Eavesdroppers and Spatially Correlated Channel Uncertainties," *IEEE Transactions on Communications*, vol. 72, no. 3, pp. 1532–1547, 2024.

[71] C. Liu, Y. Li, D. W. K. Ng, J. Yuan, and L. Dong, "Dual-Sided Active-IOS-Enhanced Secure Multi-Cell Systems Exploiting Eavesdroppers' Statistical CSI," *IEEE Transactions on Wireless Communications*, vol. 24, pp. 2663–2679, 2025.

[72] S. Hu, Z. Wei, Y. Cai, C. Liu, D. W. K. Ng, and J. Yuan, "Robust and Secure Sum-Rate Maximization for Multiuser MISO Downlink Systems with Self-sustainable IRS," *IEEE Transactions on Communications*, vol. 69, no. 10, pp. 7032–7049, 2021.

[73] X. Yu, D. Xu, Y. Sun, D. W. K. Ng, and R. Schober, "Robust and Secure Wireless Communications via Intelligent Reflecting Surfaces," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 11, pp. 2637–2652, 2020.

[74] T. A. Le, D. W. K. Ng, and X.-S. Yang, "A rank-one optimization framework and its applications to transmit beamforming," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 1, pp. 620–635, 2024.

[75] Y. Jiang, Q. Wu, H. Hui, W. Chen, and D. W. K. Ng, "Low-altitude UAV tracking via sensing-assisted predictive beamforming," *IEEE Transactions on Wireless Communications*, vol. 25, pp. 17 975–17 988, 2026.