THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

eVTOL Aircraft Energy Overhead Estimation under Conflict Resolution in High-Density Airspaces

ICNS 2026

Alex Zongo • Peng Wei

Intelligent Aerospace Systems Lab (IASL)
Department of Mechanical and Aerospace Engineering
George Washington University

IASL

The Promise of Advanced Air Mobility

Urban Air Taxis

On-demand or scheduled aviation services transforming urban transportation

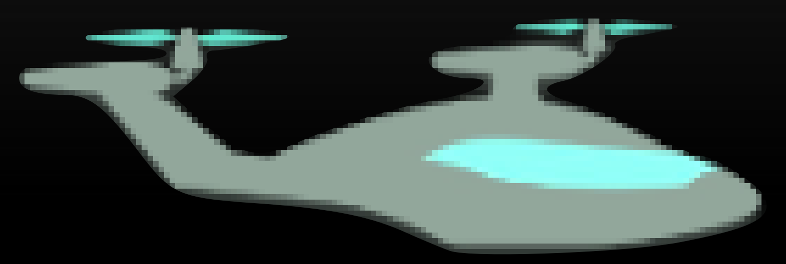

Electric Propulsion

Reduced emissions and acoustic signatures for community acceptance

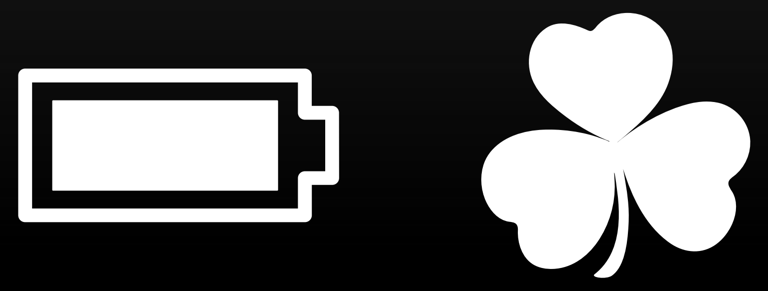

Autonomous Separation

At scale, with dozens of aircraft sharing the same airspace, tactical deconfliction is required. Aircraft should be able to maintain separation.

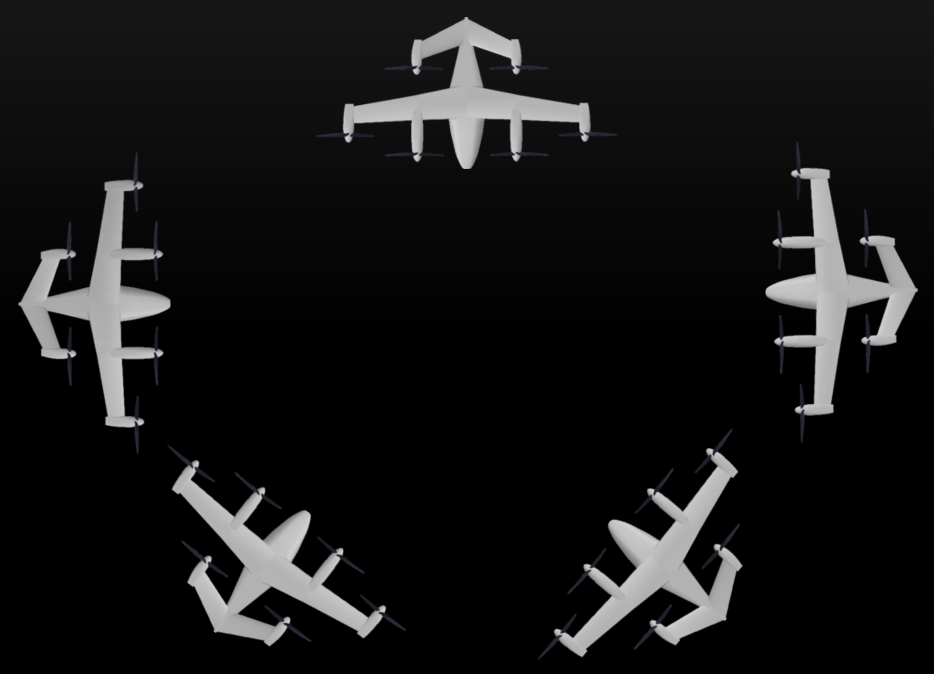

The Promise of Advanced Air Mobility

Urban Air Taxis

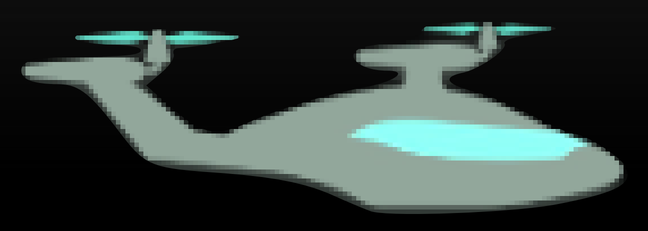

Electric Propulsion

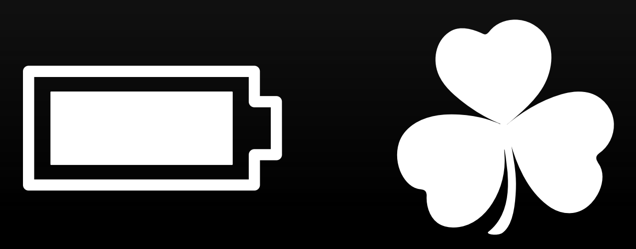

Autonomous Separation

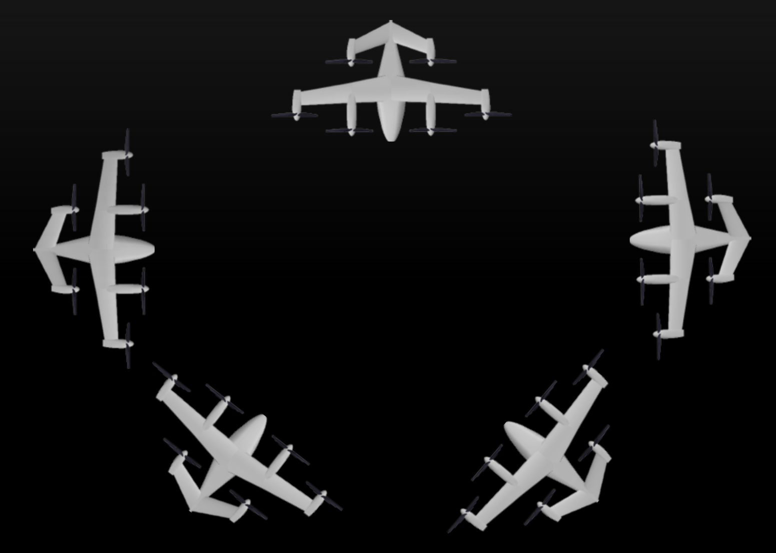

But as this vision approaches reality, a critical question emerges.

How much energy do conflict-resolution maneuvers actually cost?

And can we predict this cost before flight?

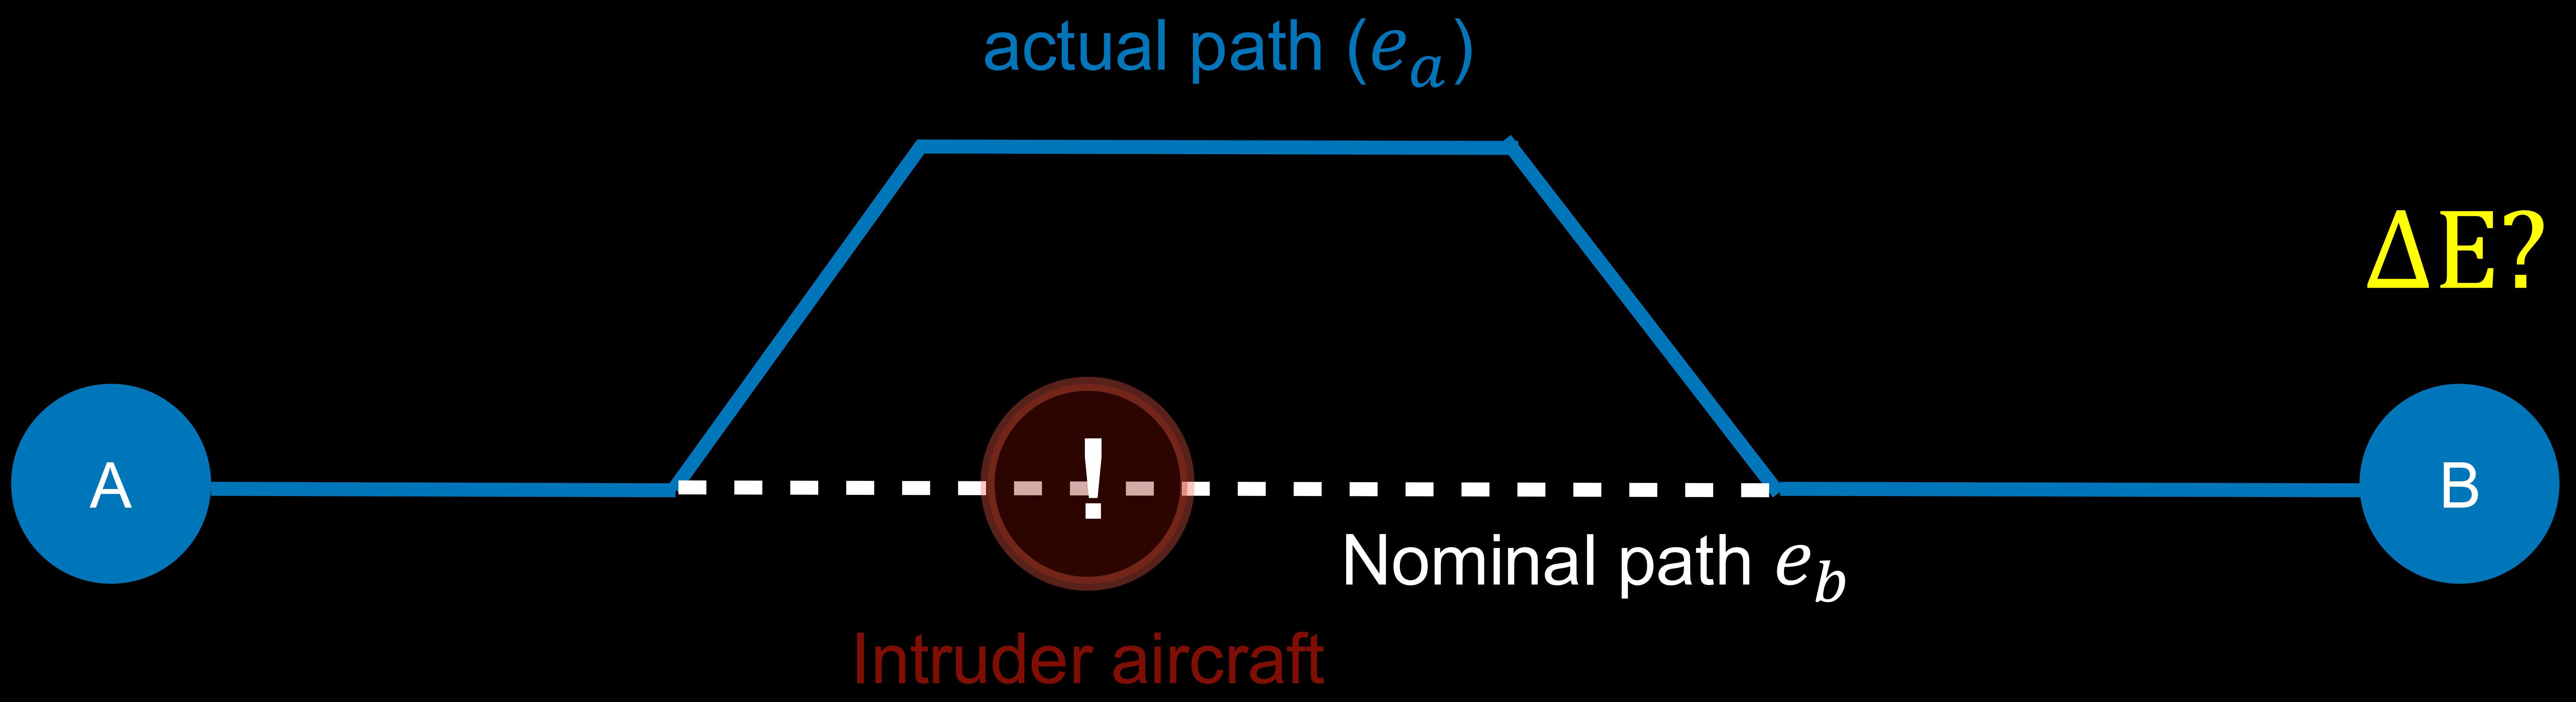

Every deviation from the nominal path consumes additional energy, but by how much?

The Energy Challenge

Battery Constraints

Energy density limitations of lithium-ion cells leave little margin for inefficiency.

Precise energy planning is essential for mission viability.

Traditional flight planning assumes nominal trajectories — there is little guidance on the cost of deconfliction.

The Operator's Dilemma

Insufficient reserves

Safety risk if conflicts require more energy

VS

Excessive conservatism

Increases charging time and slows down operations

We need quantitative characterization of deconfliction energy costs.

The Research Gap

Three research areas have advanced independently:

eVTOL Energy Models

Physics-based power models, mission planning, battery dynamics

Conflict Resolution

Modified Voltage Potential (MVP) algorithm, Reinforcement Learning, optimization methods

Predictive Machine Learning

Power estimation, traffic prediction, but less emphasis on uncertainty quantification

The Gap:

While energy models and conflict resolution have advanced independently, their intersection remains unexplored.

Can we quantify how conflict-resolution maneuvers affect eVTOL energy consumption and/or provide uncertainty-aware predictions?

Our Contributions

1

Integrated Framework

We integrate an eVTOL power model with a tactical deconfliction algorithm (MVP) in a traffic-level simulation (BlueSky [1]).

2

Empirical Characterization

We empirically show that MVP is energy-efficient: with low energy overhead in high-density air traffic.

3

Uncertainty Aware Prediction

We present a machine learning (ML) model that estimates energy consumption with calibrated confidence bounds, pre-flight (before flying a cruise segment).

[1] J. M. Hoekstra and J. Ellerbroek, “BlueSky ATC simulator project: An open data and open source approach,” in Proc. 7th Int. Conf. Research in Air Transportation (ICRAT), 2016.

Research Methodology

eVTOL Power Model → MVP Deconfliction → Traffic Simulation → Machine Learning Prediction

Key Parameters:

- Sector radius: 10 nm
- Cruise altitude: 2,000 ft
- Traffic density: 10-60 aircraft

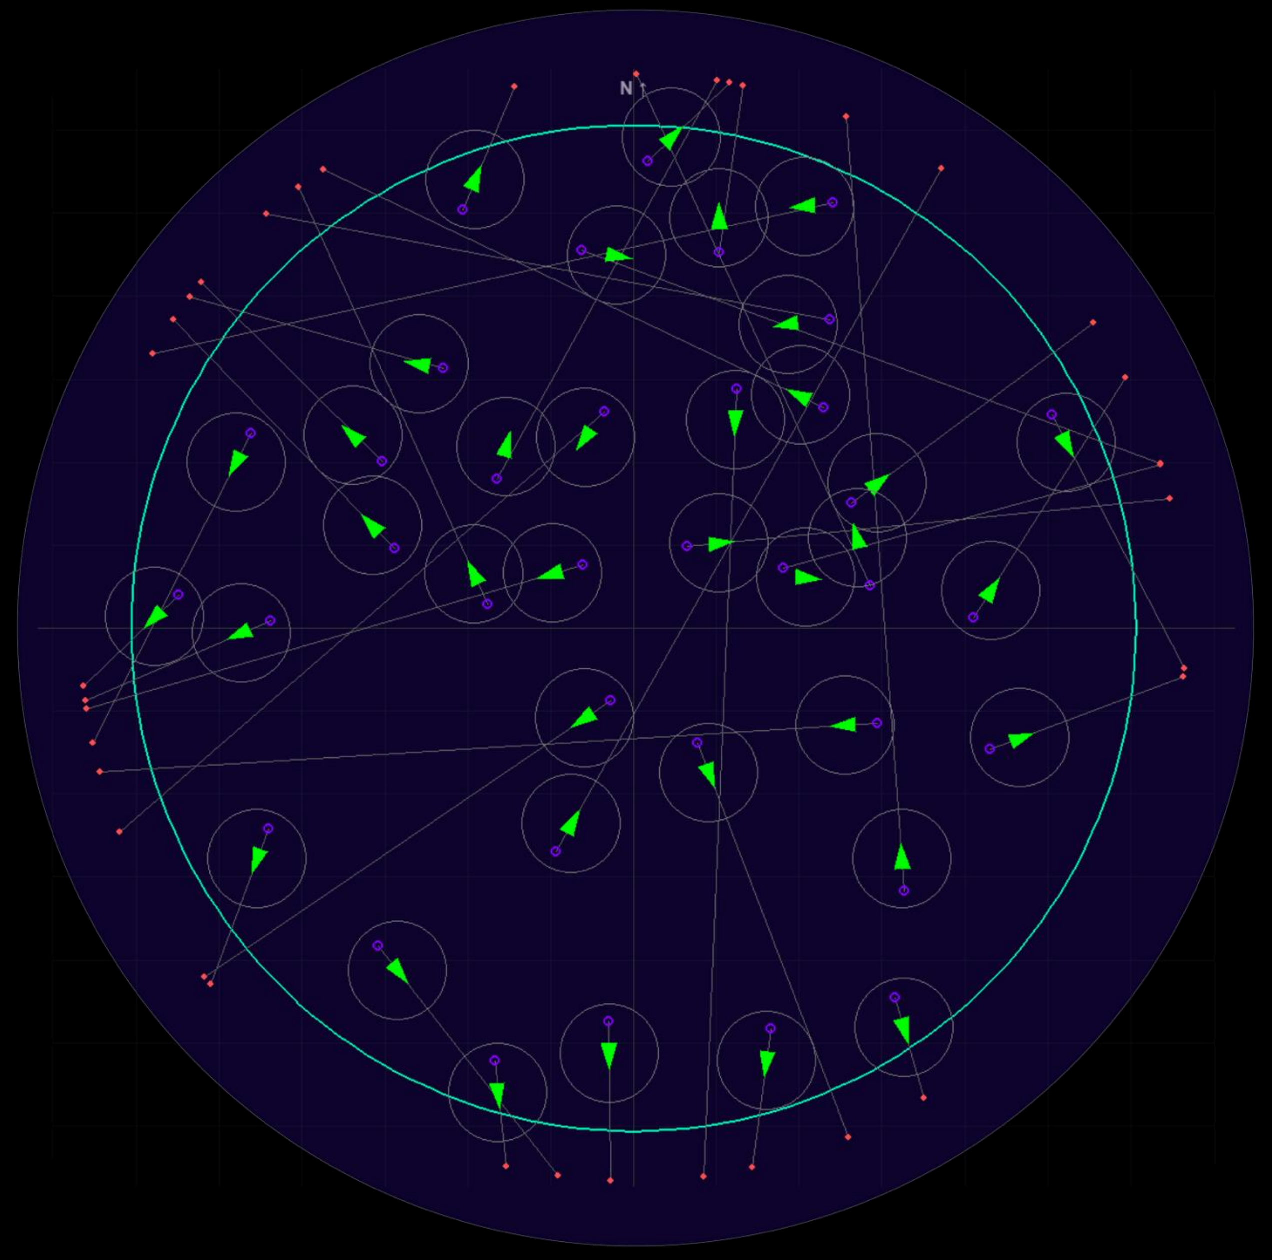

Radar snapshot of traffic in a 10 nm sector radius where the circles around each aircraft (in green) denote separation buffers; the trails show the flight plans.

eVOL Power Consumption Model

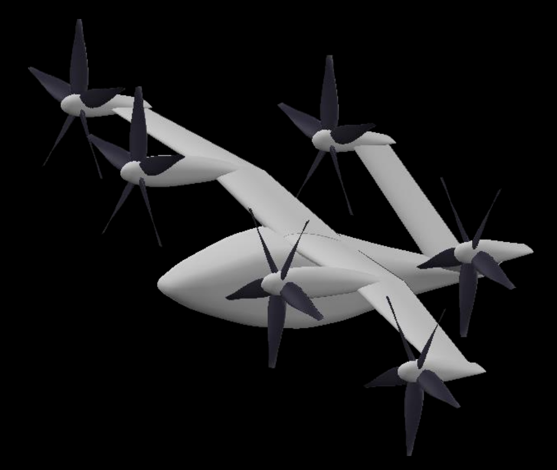

Joby S4-class tilt-rotor configuration [2]

Parameter	Value
Weight (lbs)	4,000
Number of Passengers	4
Wing Span (ft)	35
Cruise Speed (kt)	~157
Number of Rotors	6
Rotor Diameter (ft)	9.5
Max Shaft Power (kW)	6 × 115
Battery Capacity (kWh)	136

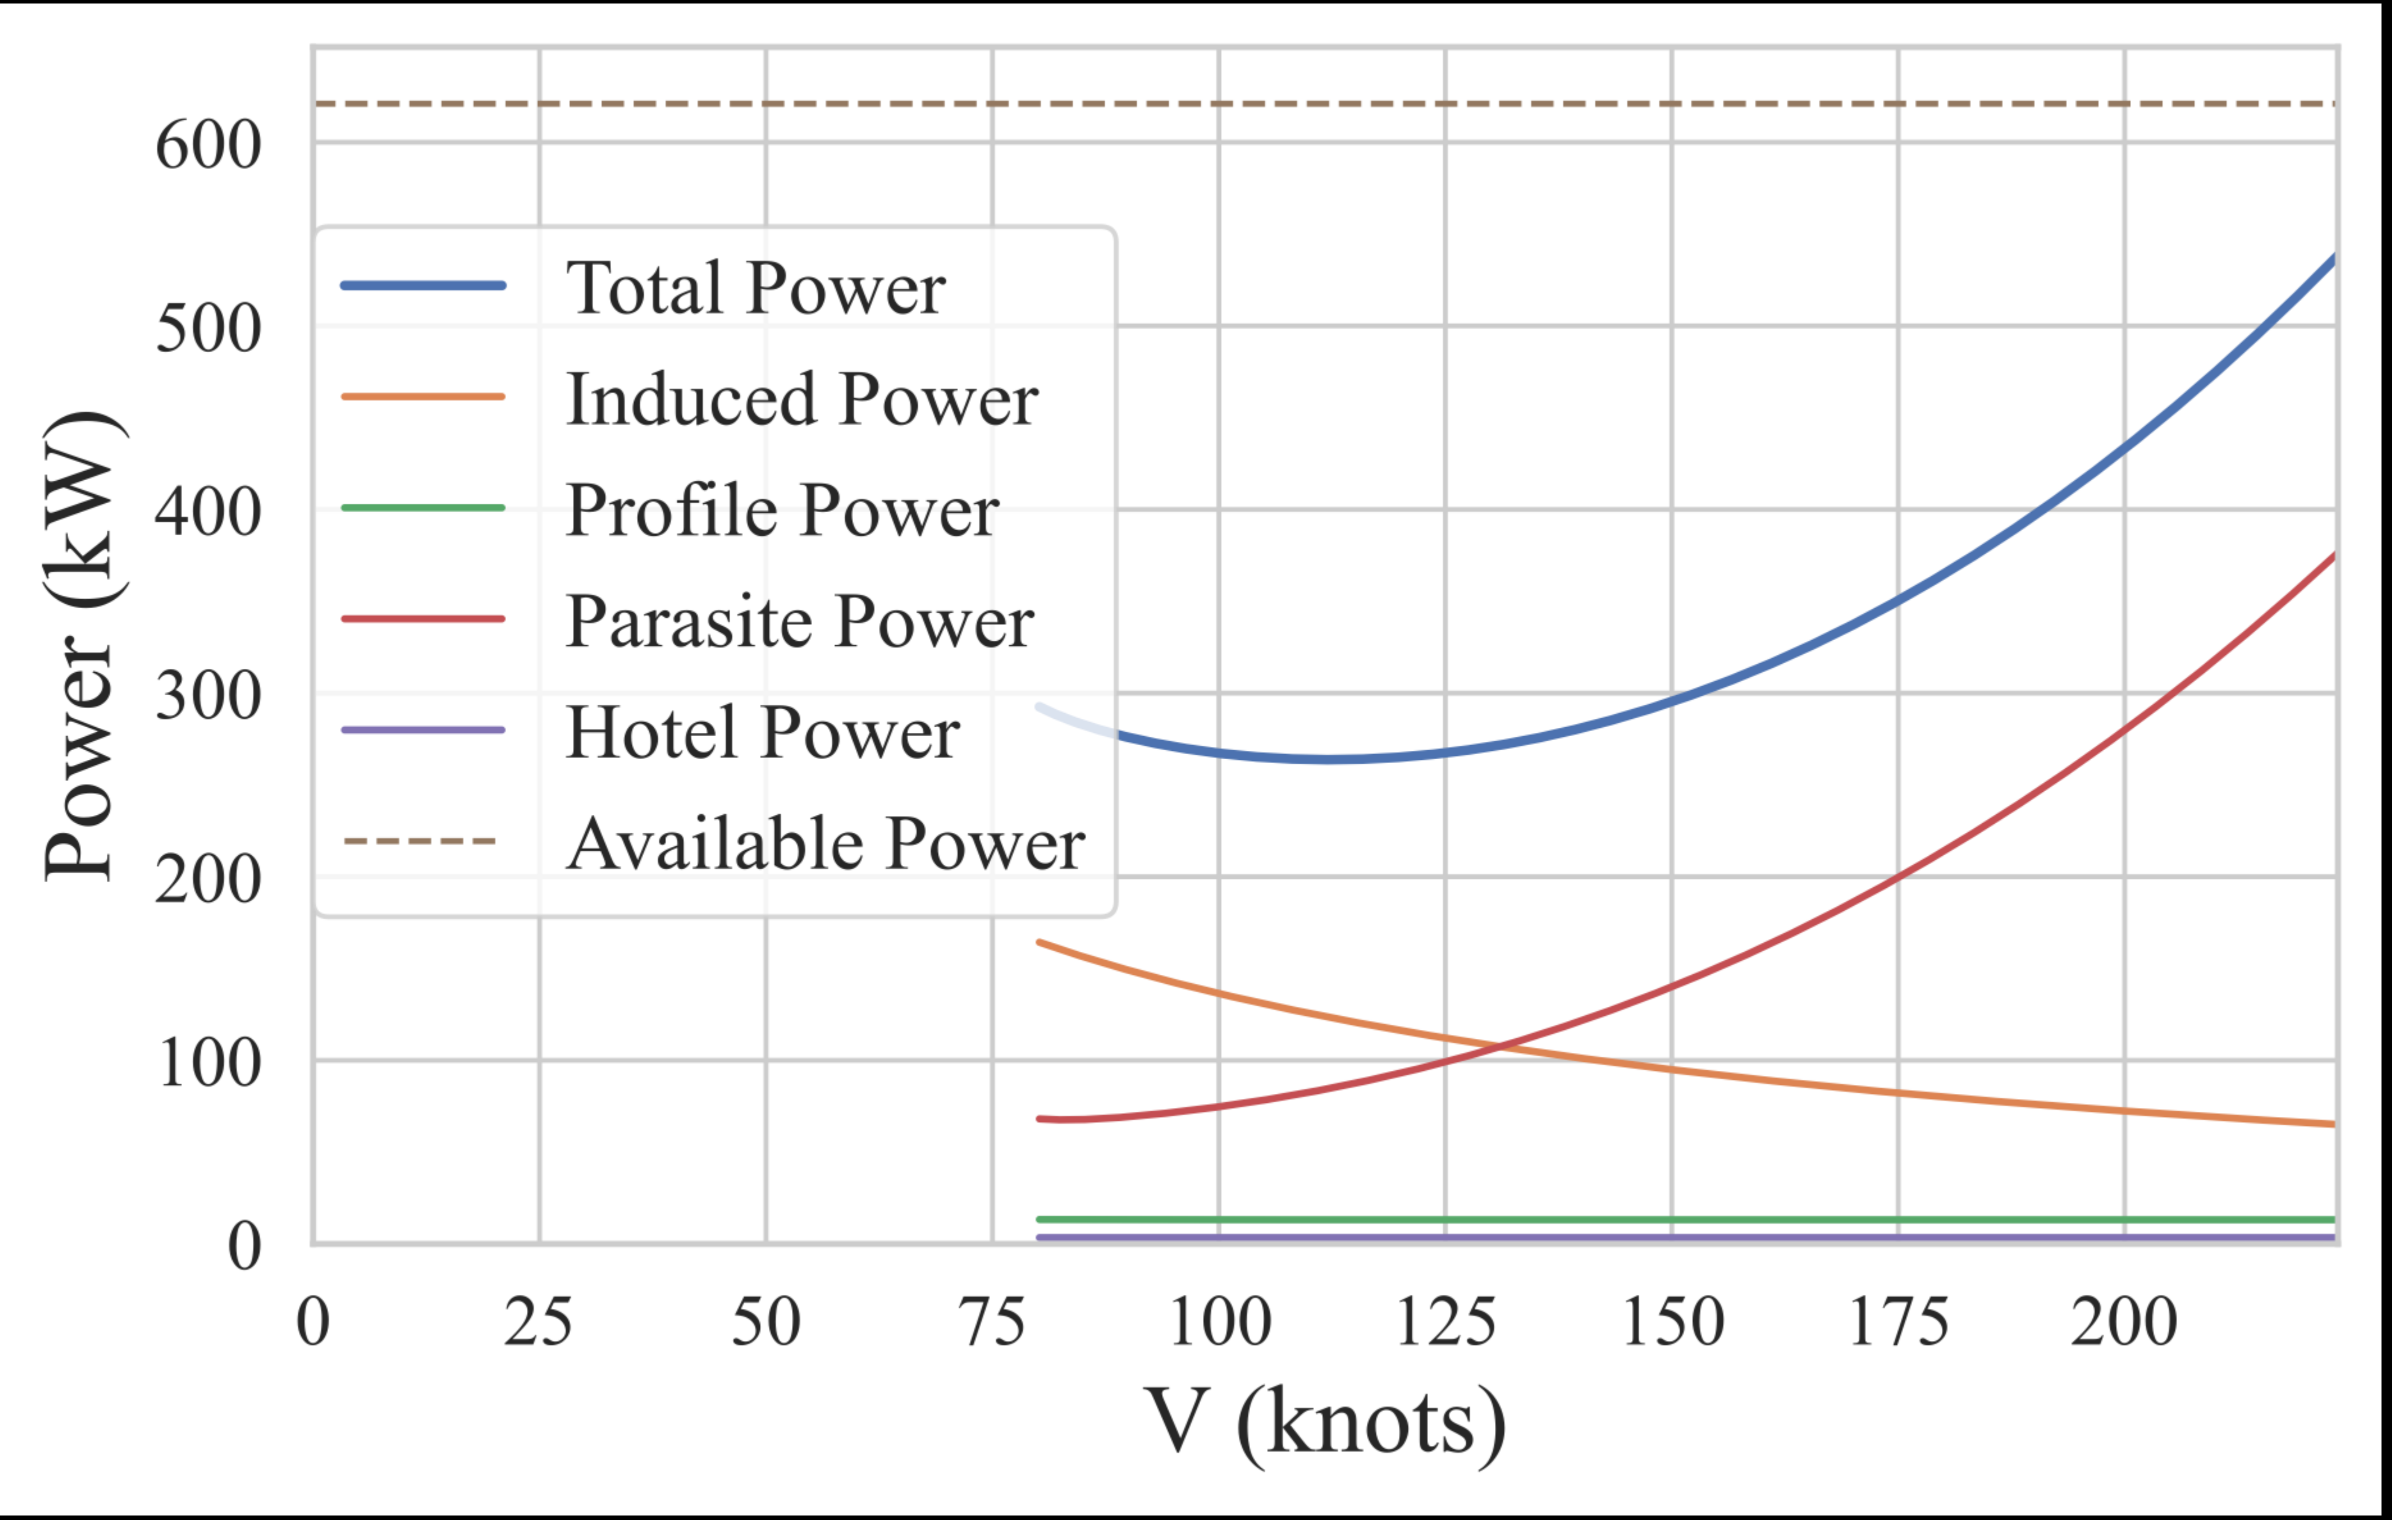

Power velocity relationship shows typical tilt-rotor behavior with substantial margin throughout cruise envelope ($\geq$ 85kt)

Total Power Components (Cruise)

$$P_{\text{total}} = (P_{\text{induced}} + P_{\text{profile}} + P_{\text{parasite}})/\eta_{drv} + P_{\text{hotel}}$$

[2] V. R. Gonzalez and J. L. Huynh, “Integrating Aircraft Performance in Traffic Flow Management Analysis for Advanced Air Mobility,” in AIAA AVIATION FORUM AND ASCEND 2025, 2025.

MVP Tactical Deconfliction

Modified Voltage Potential algorithm (Eby 1995, Hoekstra et al. 2002)

How MVP works

1. Detect conflict when the predicted closest point of approach (CPA) < protected zone radius within the look-ahead time
2. Compute minimum velocity change to move CPA to protected-zone boundary
3. Both aircraft maneuvers are complementary without explicit communication
4. Recover to original velocity when conflict is resolved

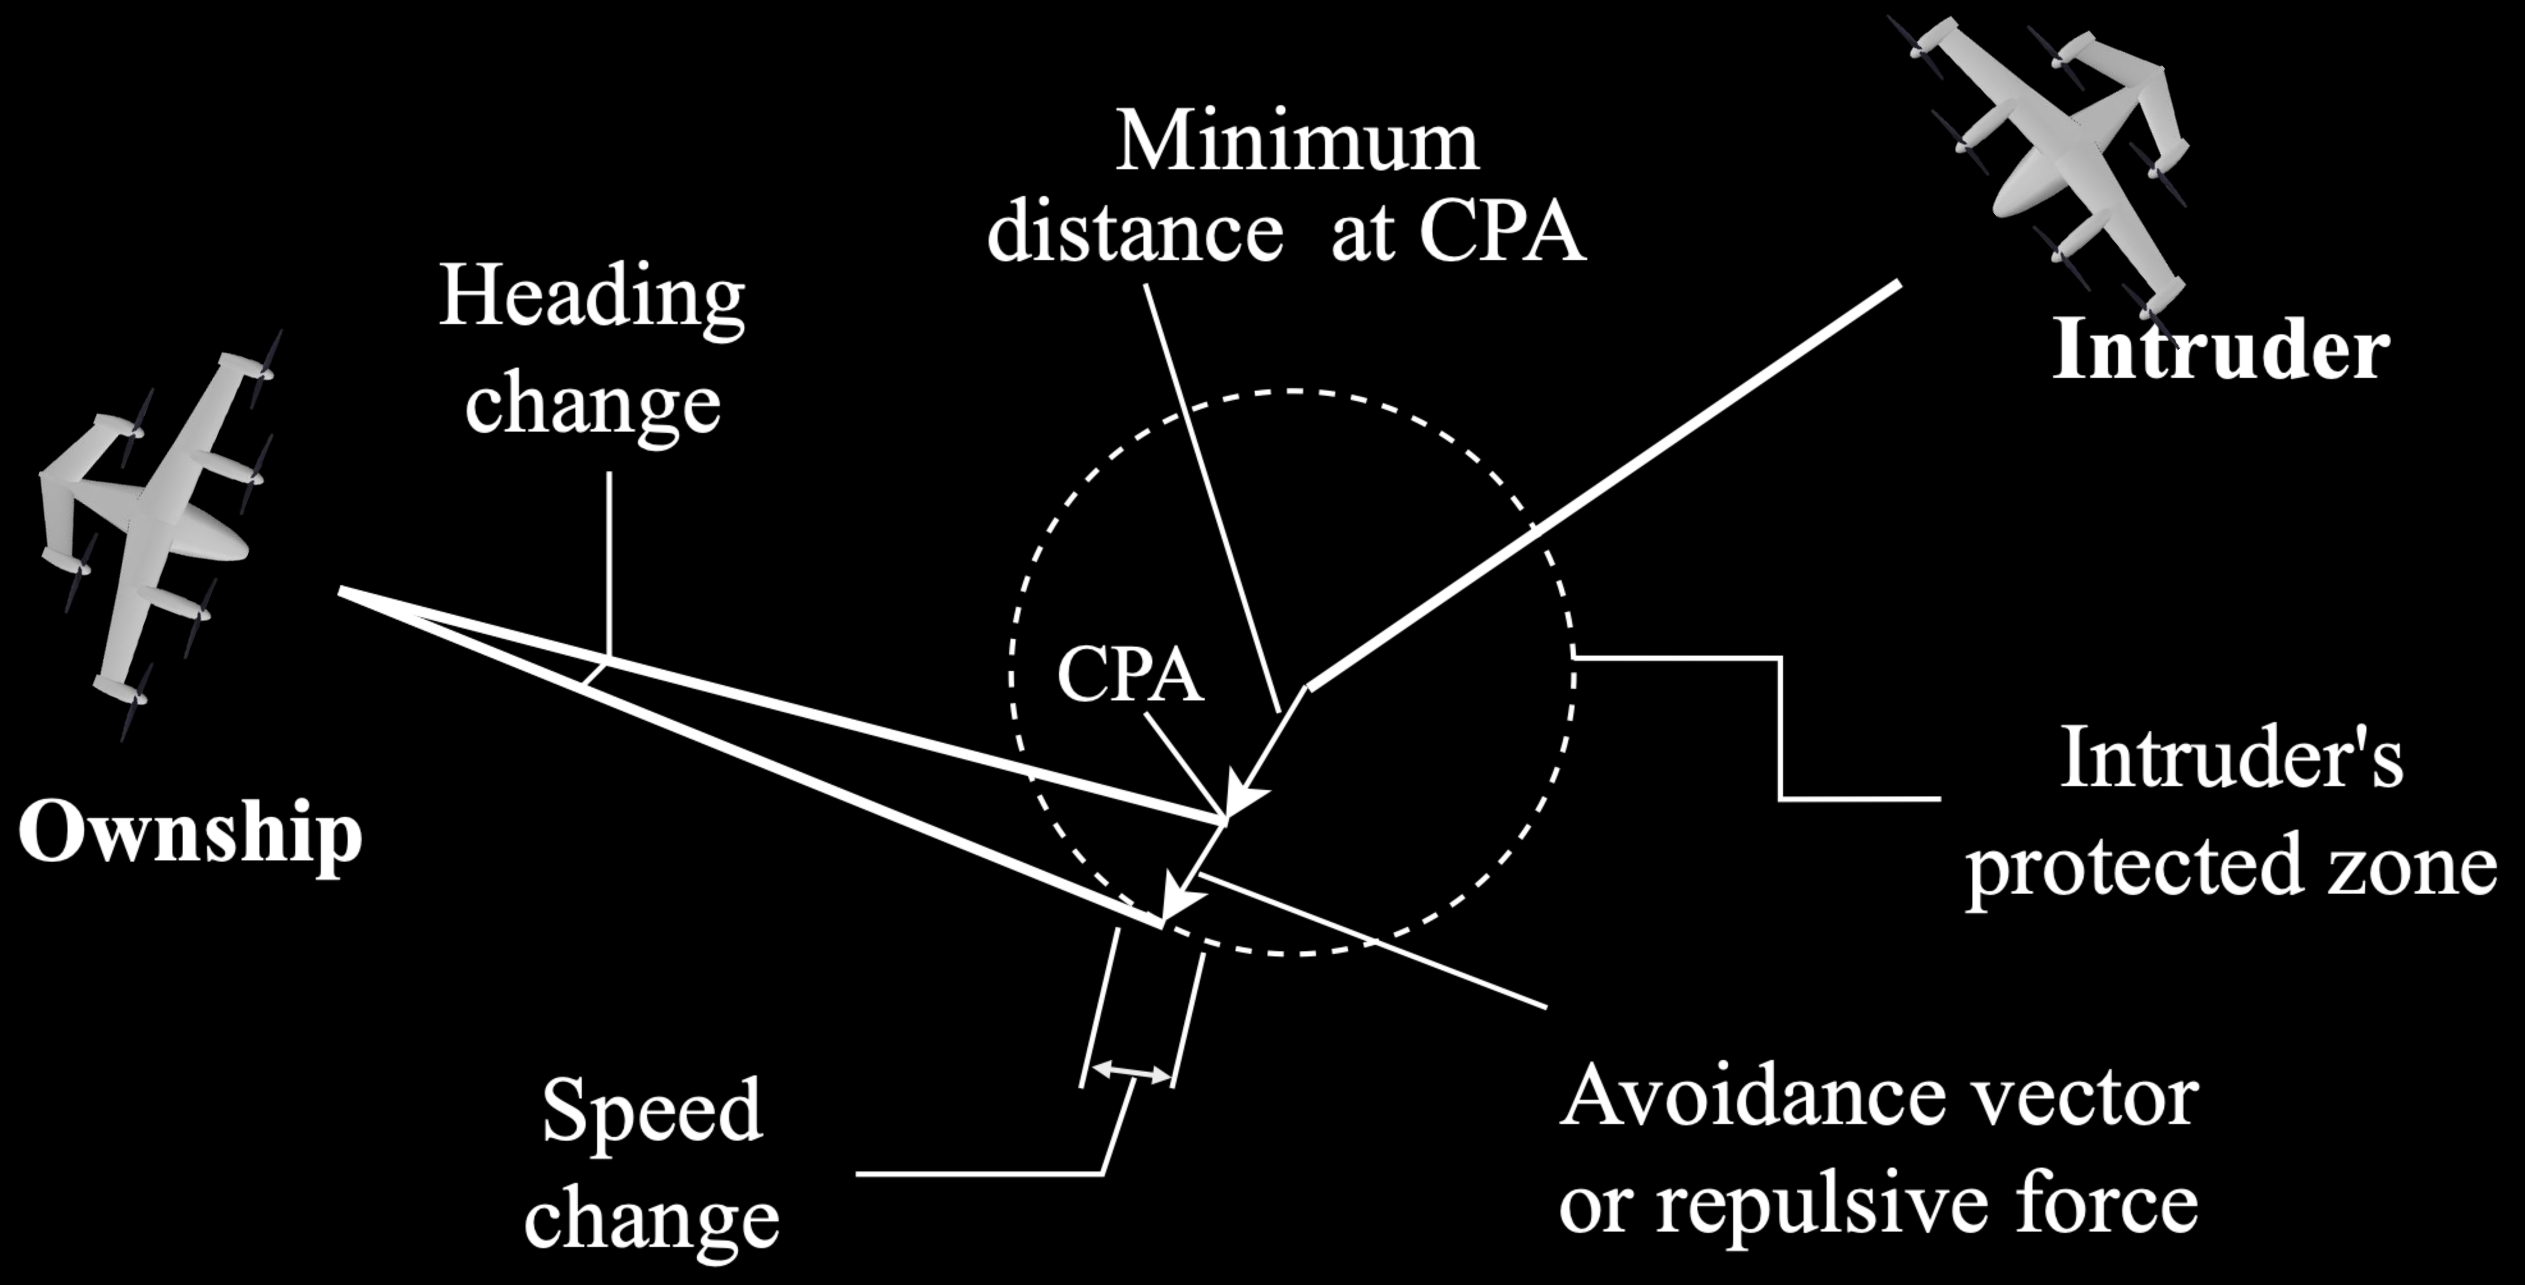

MVP Tactical Deconfliction

Modified Voltage Potential algorithm (Eby 1995, Hoekstra et al. 2002)

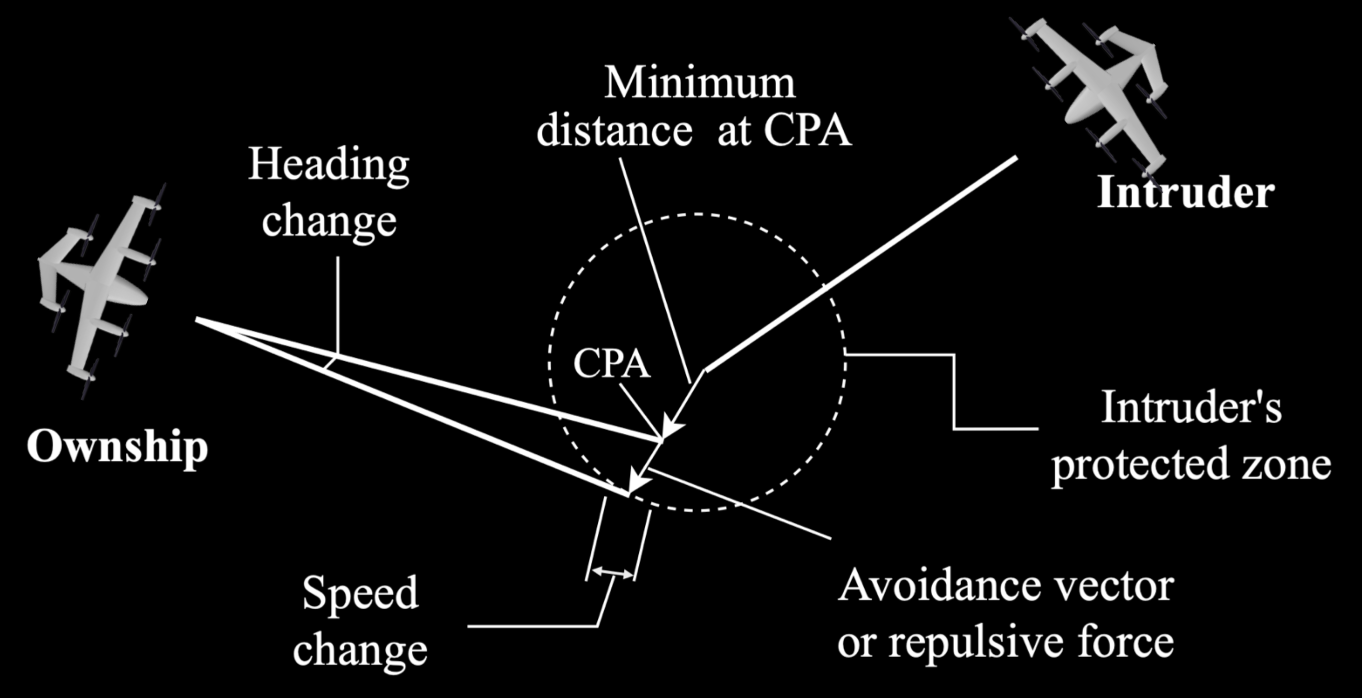

Why MVP?

- Minimal path deviation
- Deterministic and explainable
- Decentralized: no central coordination required
- Widely studied baseline in the literature
- Certification-friendly: auditable decision logic

Key Parameters:

Protected zone:	0.6 nm
Lookahead time:	90 s
Speed range:	85-185 kt

Simulation Environment

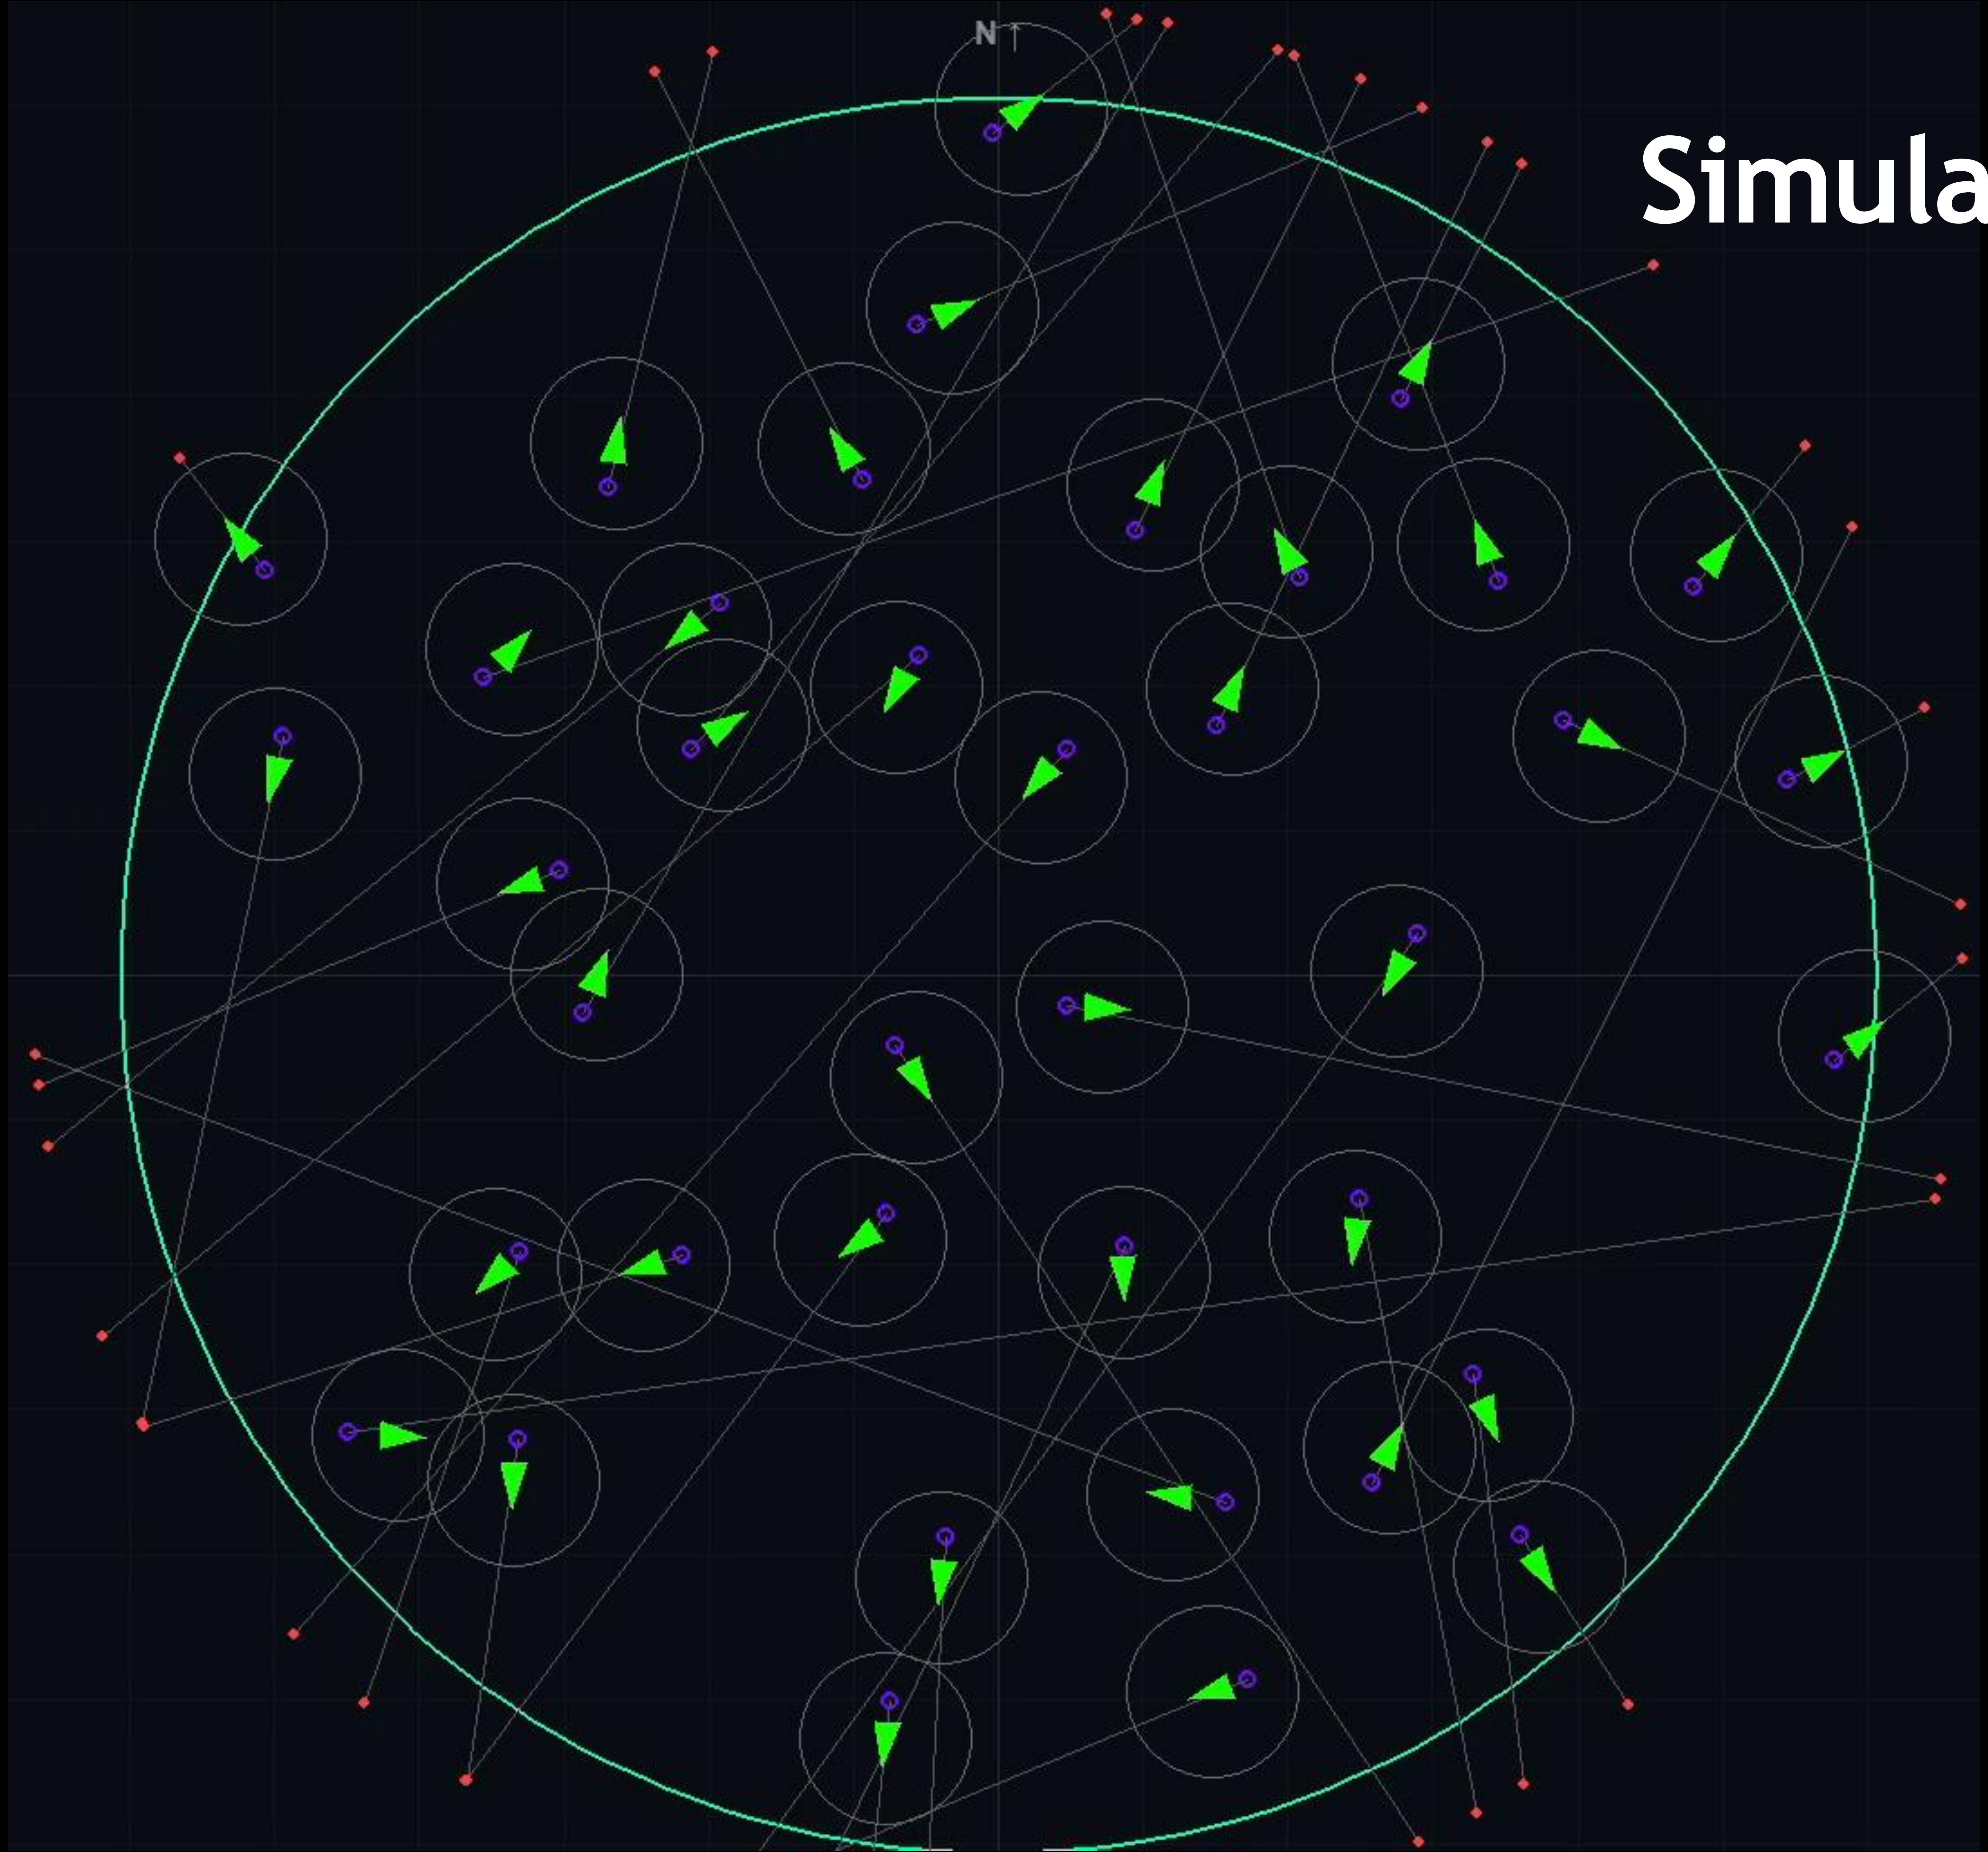

10nm radius sector

How We Measure Energy Overhead

Energy Overhead ($\delta\mathrm{E}$)

Baseline energy consumption: e_b

(conflict-free trajectory at max-range speed)

Actual energy consumption: e_a

(realized trajectory with MVP deconfliction)

$$\delta\mathrm{E} = \frac{e_a - e_b}{e_b}$$

Expressed as percentage increase relative to conflict-free baseline

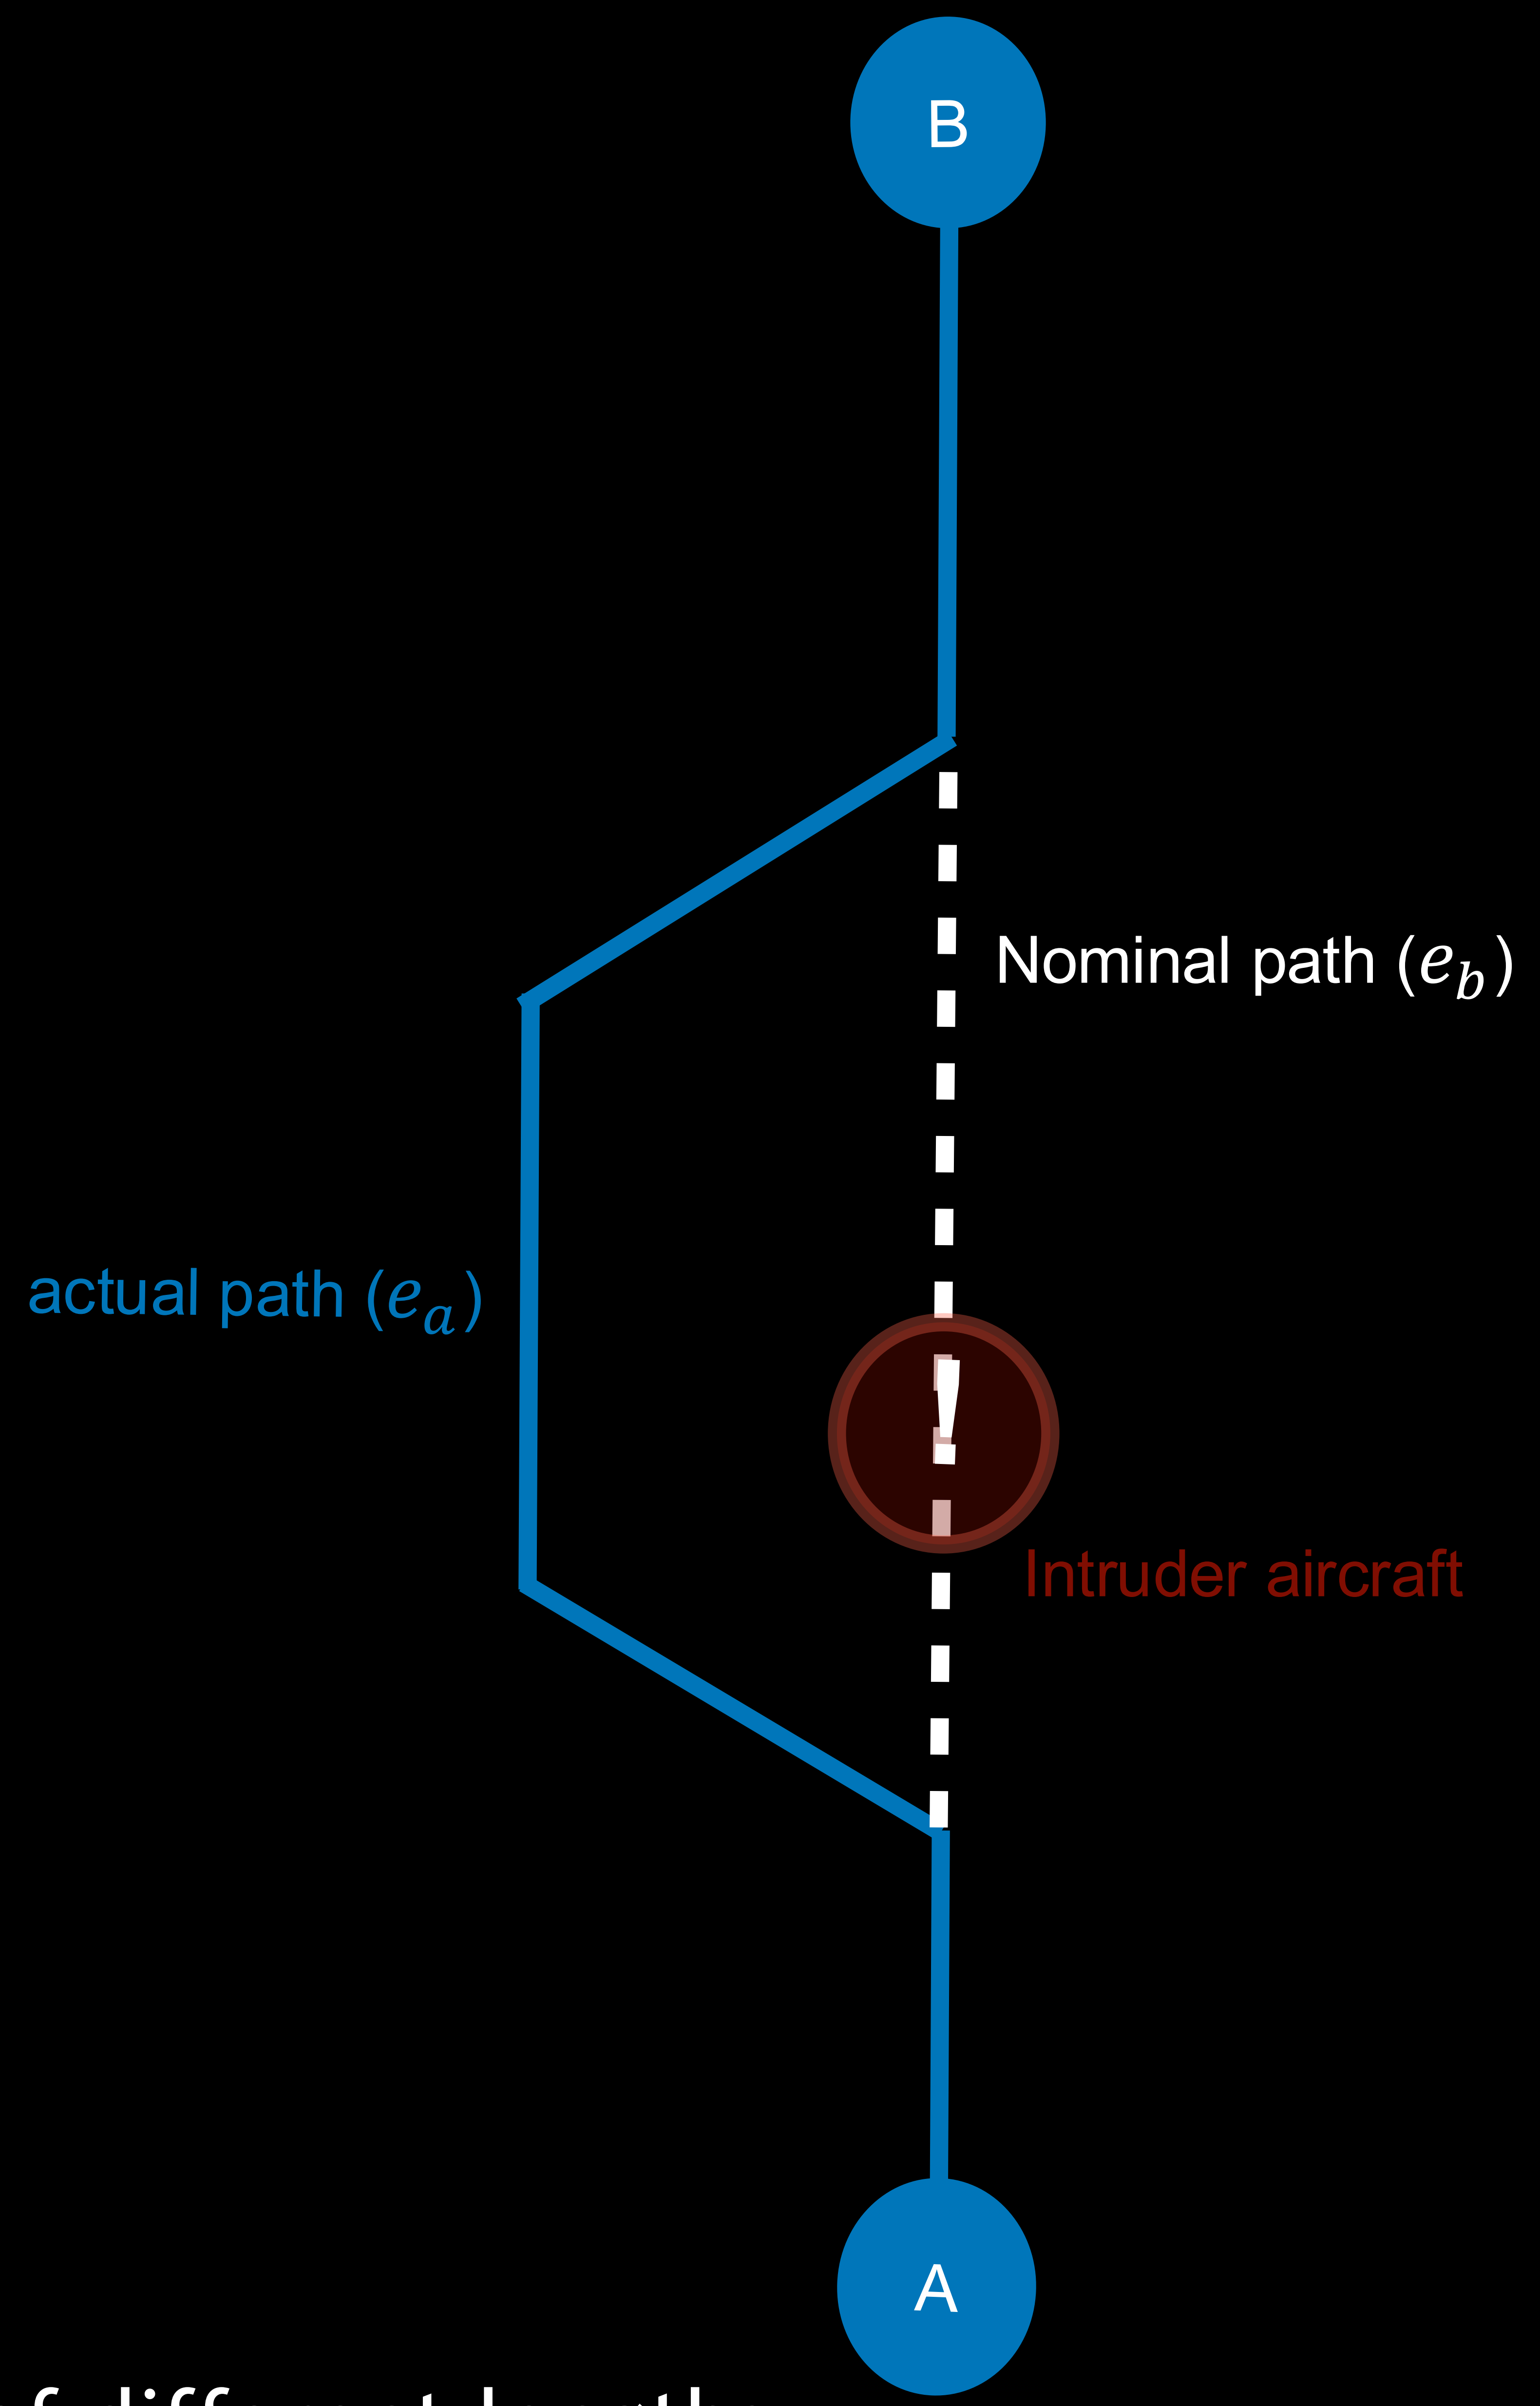

This ratio-based metric enables comparison across routes of different lengths.

Key Finding

MVP-Based Deconfliction is Energy-Efficient.

< 1.5 %

Median energy overhead across all density levels

Zero near mid-air collisions across ~70,000 simulated flights

What This Means

- The majority of transits incur negligible energy penalty.
- MVP's minimal velocity adjustments translate to minimal energy cost.
- Conflict resolution doesn't require significant reserves for typical operations.

But the Distribution Tells a Richer Story

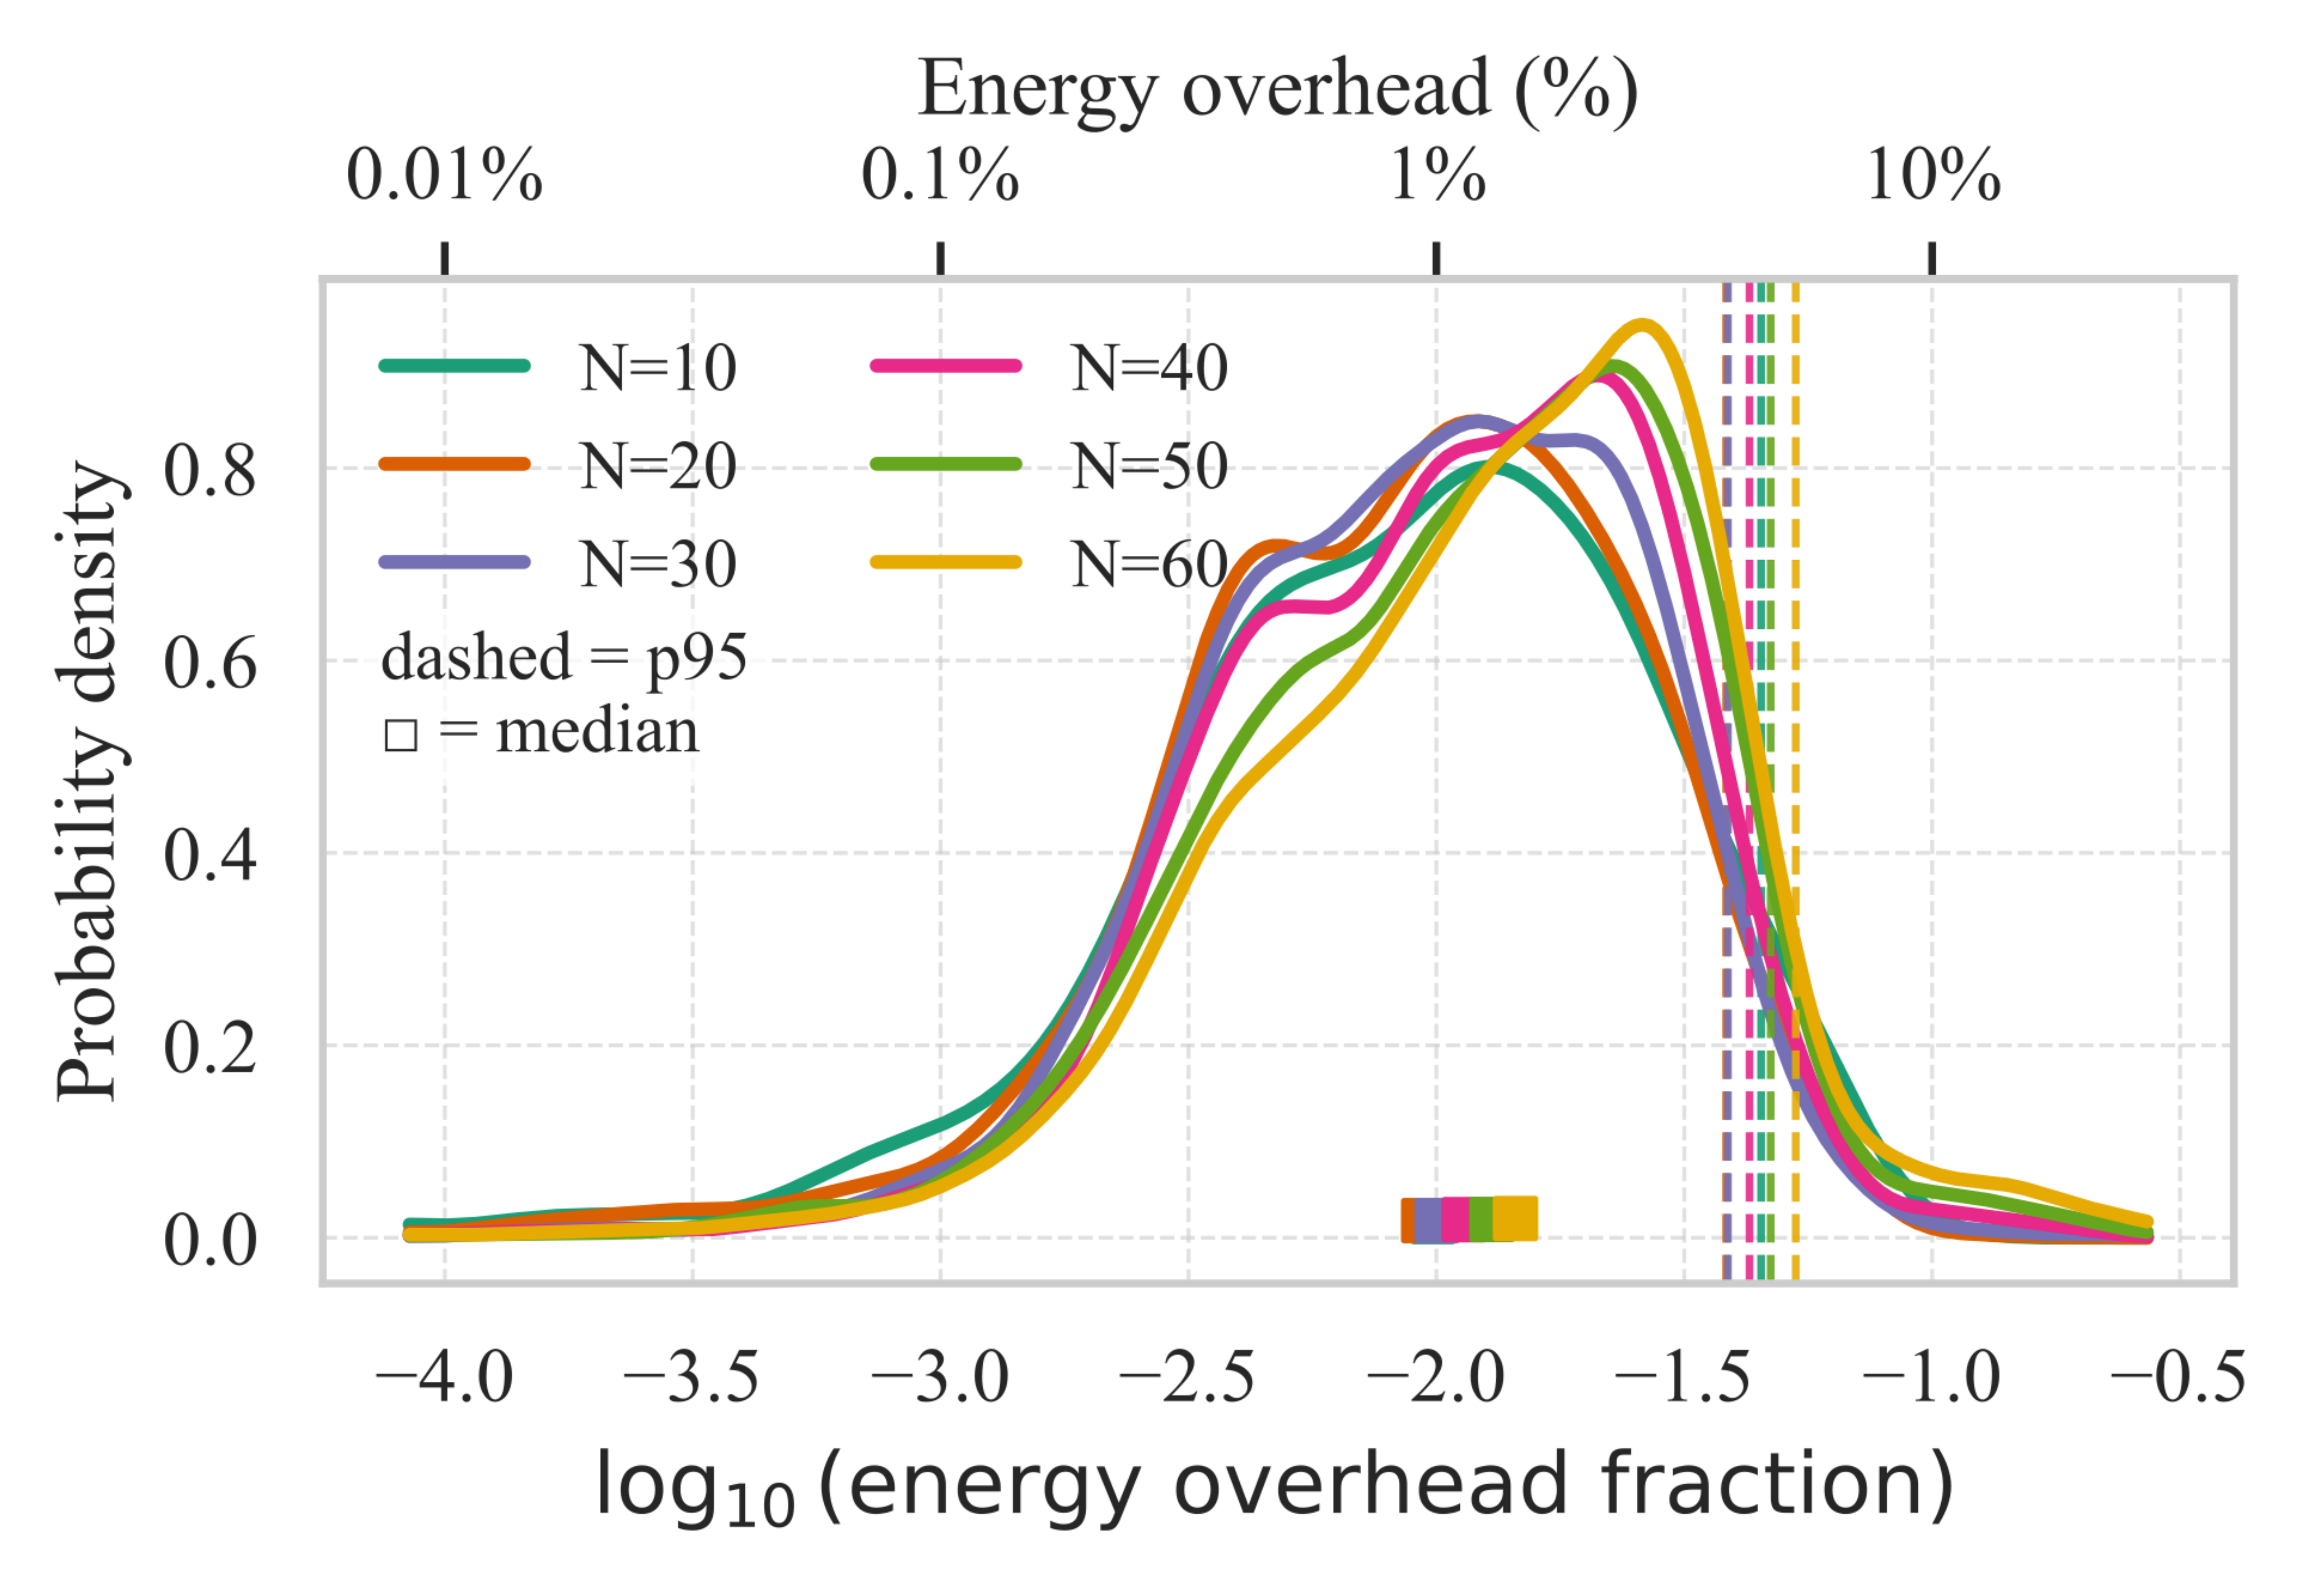

Log-scale reveals right-skewed distribution structure

Density Effect on Tails

N	P95	Max
10	4.51%	7%
30	3.87%	25%
60	5.30%	44%

Reserve Planning Insight: A 4-5% margin accommodates 95% of tactical scenarios

Why Tails Grow: The Conflict Crossover

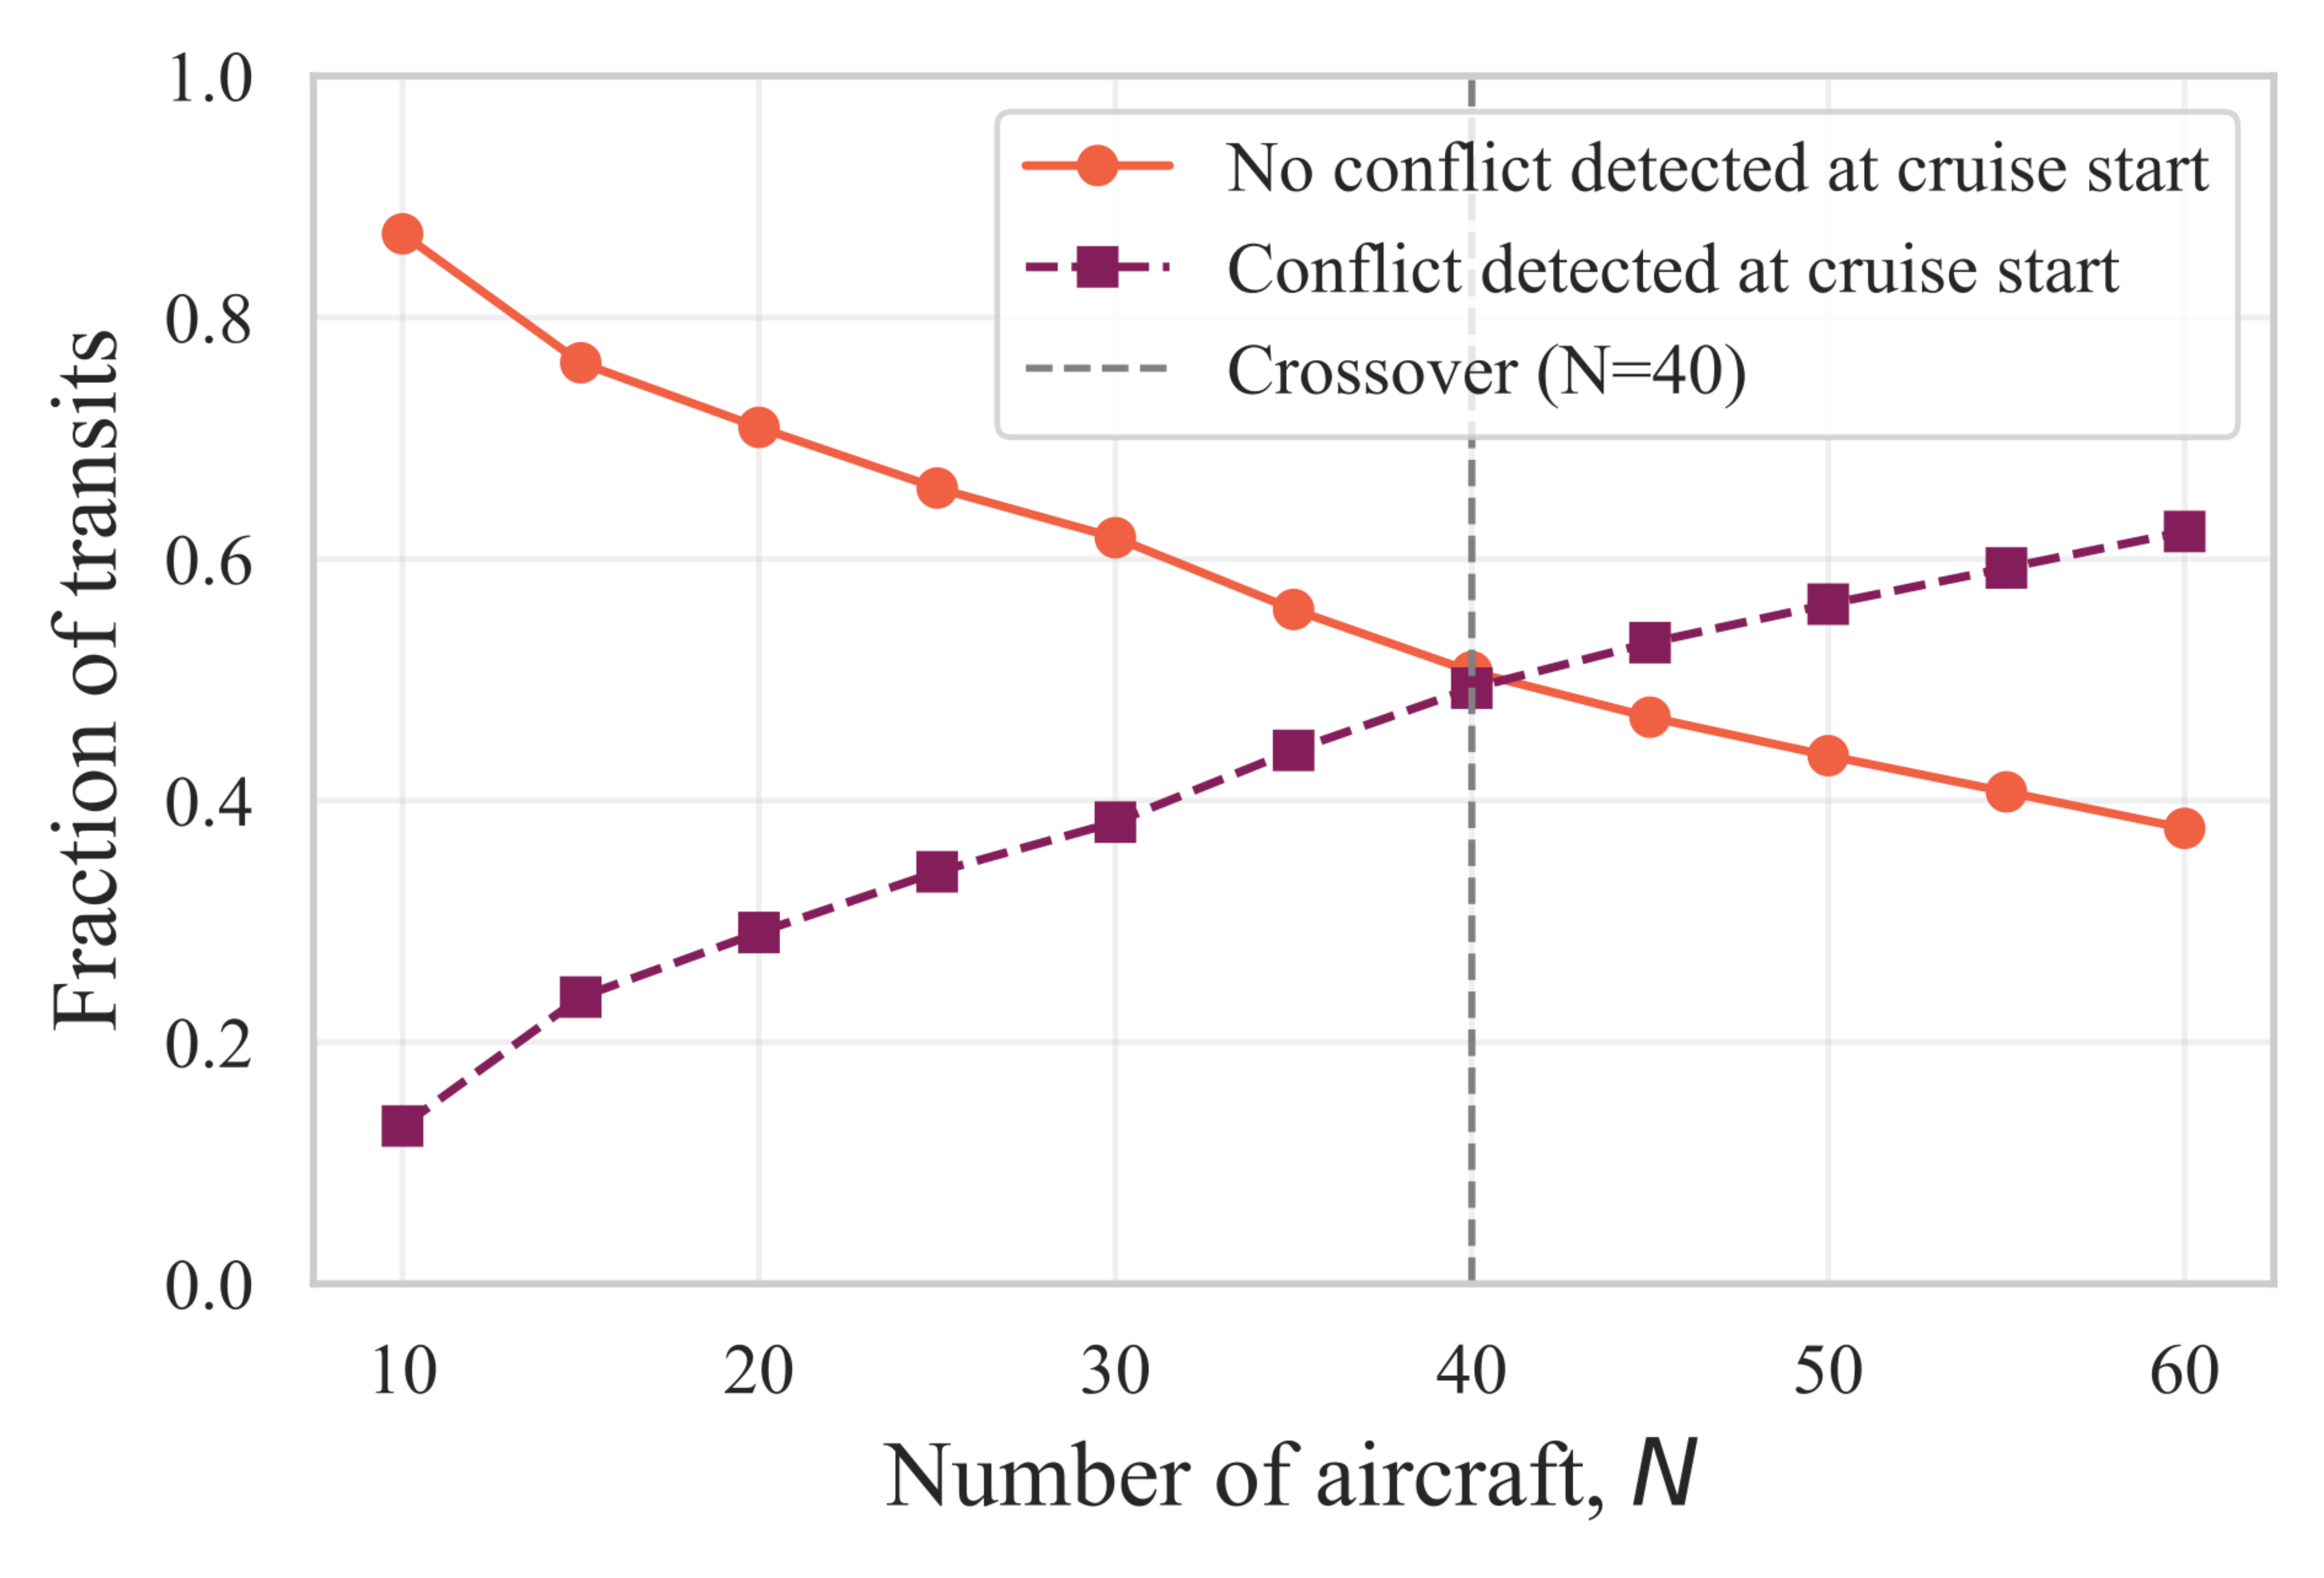

Conflict detection at cruise start

The Crossover Point

At $N \approx 40$ aircraft:

- Majority of aircraft begin cruise already in conflict.
- Multiple simultaneous conflicts become common.
- Cascading interactions drive tail energy overhead.

Below $N = 40$: 55% to 87% of flights are conflict-free at start

Above $N = 40$: Conflicts compound before resolution.

This explains why P95 grows from 4.5% to 5.3%: it's not just more conflicts, it's conflict that overlap and cascade.

From Density-based Margins to Per-flight Prediction

From Blanket to Flight-Specific Reserve Policies

Fleet-wide Percentiles

"Apply 4-5% reserve for 95% coverage"

- ✓ Simple rule
- ✓ Based on empirical data
- ✗ Same margin for all flights
- ✗ Ignores pre-flight information
- ✗ May over/under-reserve

Can we do better?

Per-flight Prediction

Query a model with this flight's conditions

- ✓ Tailored to current traffic state
- ✓ Uses initial conflict detections
- ✓ Accounts for spatial position
- ✓ Provides calibrated uncertainty bounds

Two aircraft in the same density environment but with different initial conflict states should receive different reserve recommendations.

Predicting Energy Overhead

The Problem

At mission start, we know the traffic state (density, congestion, etc) but not how conflicts will evolve:

- Which aircraft will we encounter?
- How will they maneuver?
- Will conflicts cascade?

Precise point prediction is fundamentally unreliable.

Our Approach

Instead of point estimates, provide calibrated uncertainty bounds:

- The machine learning model outputs distribution parameters
- Derive prediction intervals from the distribution
- Use conservative bounds for safety-critical reserve planning.

Key Insight: Log-transformation handles right-skewed, non-negative overhead

$z = \log(1 + \delta \mathrm{E}) \sim \mathcal{N}(\mu_z, \sigma_z^2)$ →Predict mean & variance in z-space → Invert for overhead bounds

Model Architecture

Input Features

- Speed deviation from cruise speed
- Position
- Route distance
- Conflict severity
- MVP initial resolution
- Sector congestion

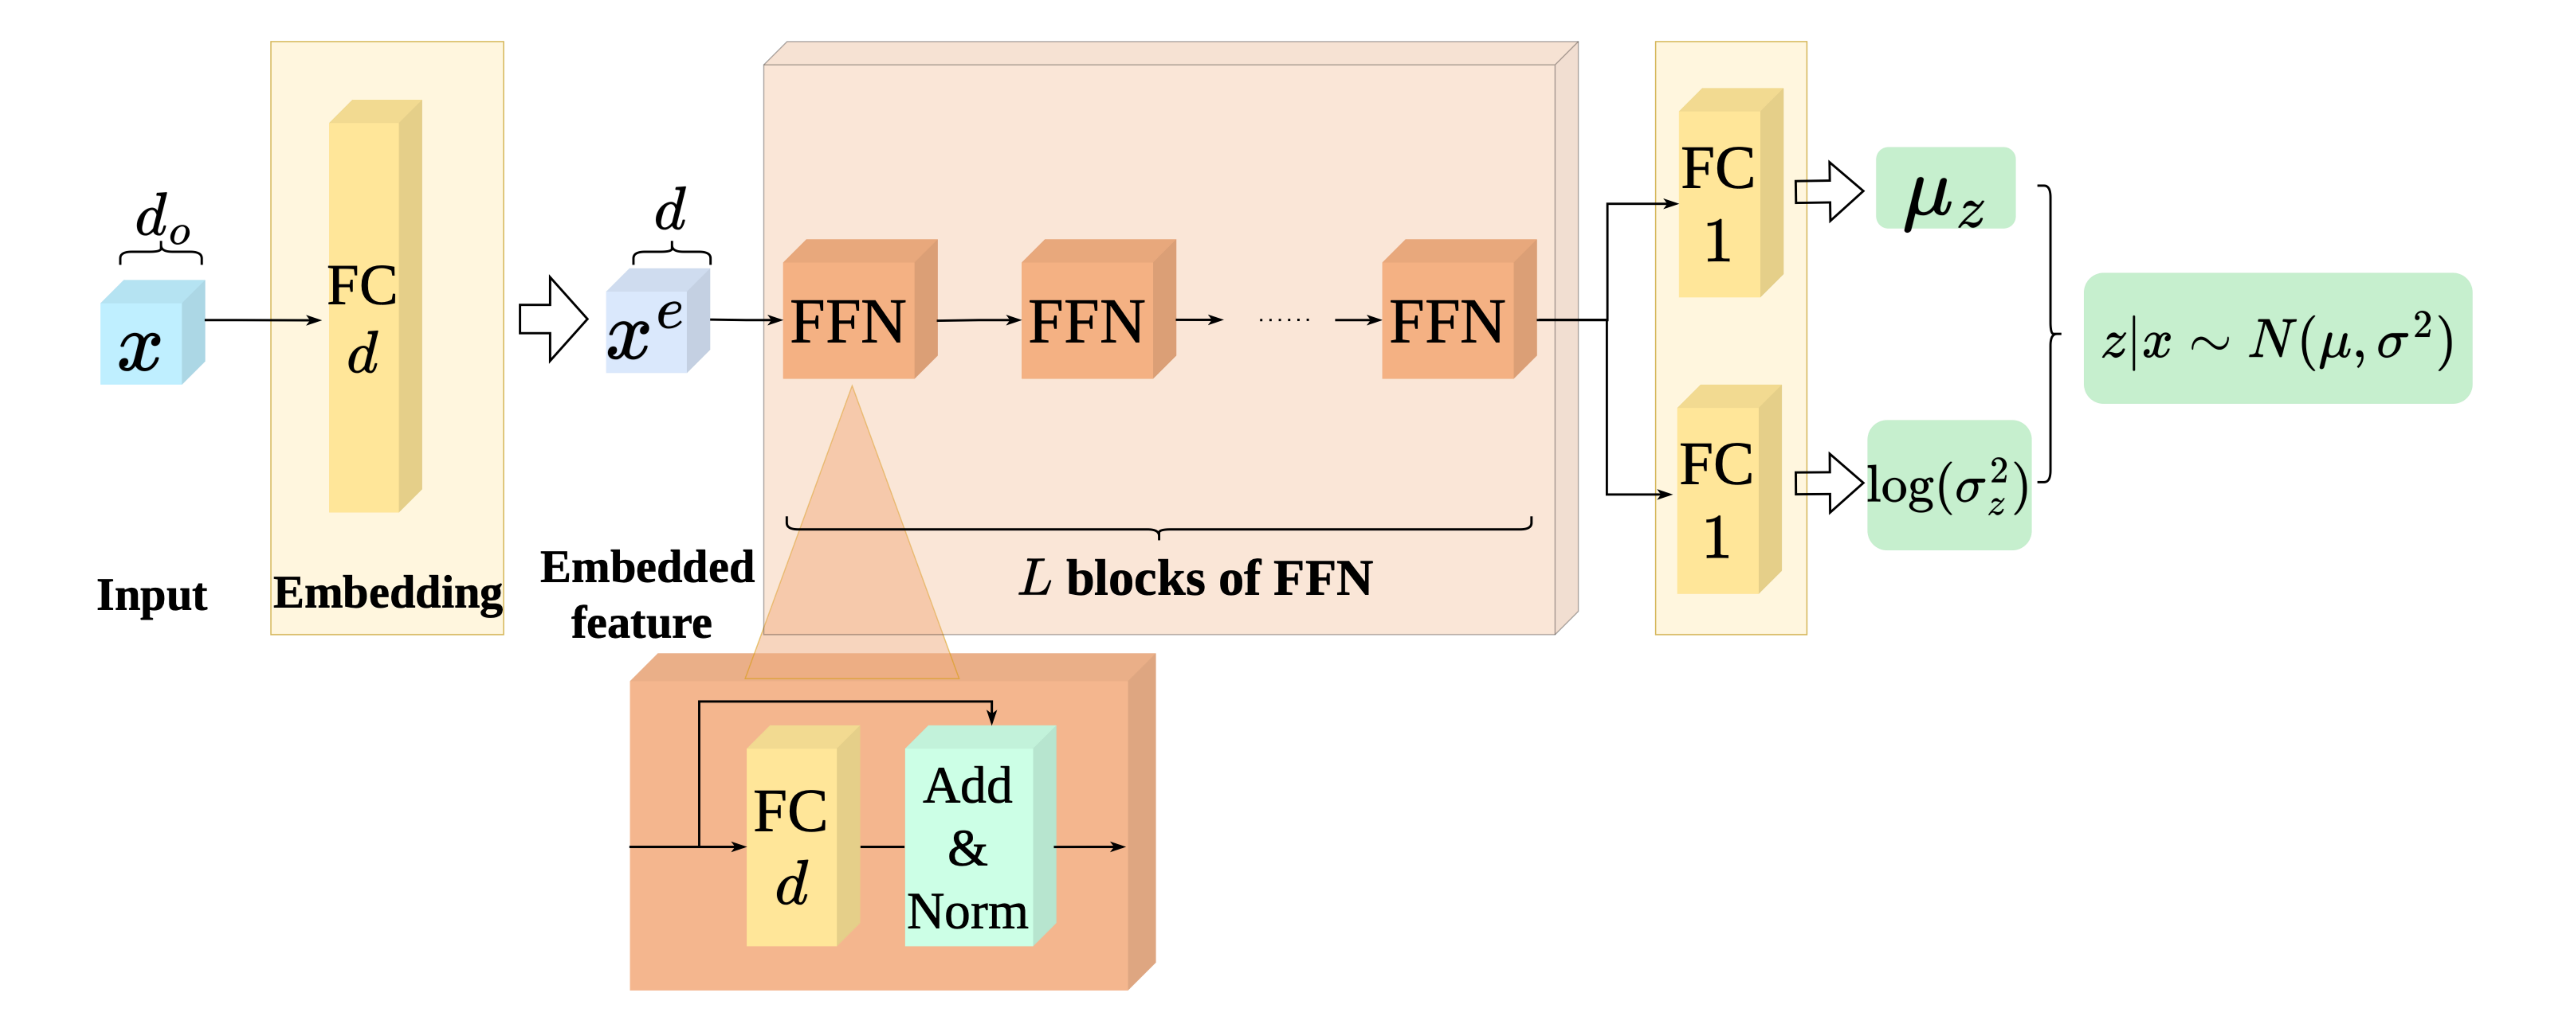

Inference

- Mean: $\mathbb{E}[\delta\mathrm{E}] = \exp(\mu_z + \sigma_z^2/2) - 1$
- Upper quantiles: $\delta\mathrm{E}_q = \exp(\mu_z + \sigma_z \cdot \Phi^{-1}(q)) - 1 \rightarrow$ reserve margin

Model Performance

Point Prediction

MAE (overall)	1.0%
RMSE	1.8%
Pearson Correlation	0.5
R^2 (1-variance)	0.22

Uncertainty Calibration ★

	Nominal	Empirical
80% interval coverage	80%	96%
90% interval coverage	90%	98%
80% interval width	—	5.3%
90% interval width	—	6.8%

Conservative Bounds for Safety-Critical Applications

Actual energy overhead estimation falls within the predicted bounds more often than the stated confidence level would suggest, which is desirable for robust planning.

MAE: Mean Absolute Error; RMSE: Root Mean Squared Error

Operational Implications

Before

Apply uniformly 5% energy reserve to all flights regardless of traffic or pre-flight conditions

→

After

Query the model with pre-flight features to receive tailored upper bounds for each flight plan

Example: an aircraft launching into:

- A low-density sector, no initial conflicts → 90% bound ≈ **3%** reserve needed
- A congested sector, multiple initial conflicts detected → 90% bound ≈ **7%** reserve needed

Limitations and Scope

Model Validation

Power model based on published parameters, not validated against flight test data.

However, relative overhead is robust to systematic biases in absolute power.

Cruise Only Scope

Simulations operate at cruise altitude thoughout.

Conflicts near vertiports during transition phases may have different characteristics.

MVP-Specific

Results apply to MVP algorithm.

Alternative methods (RL, optimization) may exhibit different energy profiles.

The ML model uses only pre-flight features; real-time updates during flight could improve accuracy.

Conclusion

1. MVP is energy-efficient.

   Autonomous tactical deconfliction may not be an energy burden.

2. Reserve planning can be tailored.

   Traffic-specific conditions enable situation-aware reserves instead of overly conservative policies.

3. ML enables uncertainty-aware prediction.

   Machine learning shows promise in estimating energy consumption under conflict resolution in high-density airspaces.

These results validate decentralized tactical deconfliction for energy-constrained eVTOL operations in high-density urban air mobility airspace.

Future Work

Extend speed envelope

Current work focused on cruise (≥ 85 kt). Plan to leverage the low-speed capabilities of eVTOL for more flexible deconfliction.

Full-envelope operations

Extend the power model to cover all flight phases: climb, descent, hovering, and cruise.

Real-time prediction

Instead of pre-flight estimation, develop a decentralized prognostic incorporating local traffic data during flight.

Thank You!

Supported by NASA Grant 80NSSC23M0059

Alex Zongo
a.zongo@gwu.edu

Intelligent Aerospace Systems Lab (IASL)

George Washington University

Backup: Simulation Parameters

Simulation Environment

Parameter	Value
Sector radius	10 nm
Cruise altitude	2,000 ft
Max-range speed (Vbr)	157 kt
Traffic count (N)	10-60
Runs per density	200
Total transits	~71,767

Conflict Detection & Resolution

Parameter	Value
Protected zone radius	0.6 nm
Look-ahead time	90 s
Permitted speed range	85-185 kt
Exit bearing range	60°–180°
Initial separation factor	$\alpha \approx 1.60$

Backup: Power Model Components (1)

Total Power Components (Cruise)

P_{ind}	Induced power (rotor thrust)	$\propto \kappa \cdot T \cdot v_i (2)$
P_{profile}	Profile power (blade drag)	$\propto \frac{C_d^b \sigma}{8}(1 + k_\mu \mu_r^2) \cdot \rho A (\Omega R)^3 \cdot N_{\mathrm{rotors}} (3)$
P_{parasite}	Parasite power (airframe drag)	$\propto D \cdot V (4)$
P_{hotel}	Hotel loads (~ 2kW)	
P_{total}	Total power	$P_{\mathrm{total}} = \frac{P_{\mathrm{ind}} + P_{\mathrm{profile}} + P_{\mathrm{parasite}}}{\eta_{\mathrm{drv}}} + P_{\mathrm{hotel}} (5)$

Backup: Power Model Components (2)

Aerodynamic Drag

We compute the aerodynamic drag coefficient of each major component of the aircraft: wing, fuselage, tail surfaces, propulsion pods, and landing gear

In cruise flight, the zero-lift drag coefficient for each component, with the exception of landing gear and propulsion pods, follows the standard form

$$C_{D,0}^{(i)} = FF^{(i)} \cdot C_f^{(i)} \cdot \frac{S_{wet}^{(i)}}{S_{ref}},$$

Where:

$FF^{(i)}$ is the form factor that incorporates thickness,

$C_f^{(i)}$ is the skin friction coefficient from the Prandtl-Schlichting turbulent flat-plate formula,

$S_{wet}^{(i)}$ is the wetted area of component i

S_{ref} is the wing area.

Shevel - Ch 11, pg 181 and Raymer Aircraft Design Ch 12 pg 283

Backup: Power Model Components (3)

Aerodynamic Drag

We compute the aerodynamic drag coefficient of each major component of the aircraft: wing, fuselage, tail surfaces, propulsion pods, and landing gear

In cruise flight, the zero-lift drag coefficient for each component, with the exception of landing gear and propulsion pods, follows the standard form

At cruise:

For the landing gear:

$$C_{D,0} \approx 1.16 \cdot (b_{gear} \cdot d_{gear})/S_{ref}$$

Hoerner Aircraft Components Section - Ch 13, pg 13-15

For each pod:

$$C_{D,0} \approx (1.1 \cdot \sin^3(\theta_{tilt} - \alpha) + 0.0204) \cdot (l_{pod} * d_{pod})/S_{ref}$$

Hoerner Elliptical Section - Ch 03, pg 3-11

Backup: Power Model Components (1)

Aerodynamic Drag

Overall, given each drag coefficient defined with respect to the same reference area S_{ref} :

$$C_{D,p}^{\mathrm{cruise}} = \sum_i C_{D,0}^{(i),\mathrm{cruise}}. \qquad D_p = C_{D,p}^{\mathrm{cruise}} \cdot q \cdot S_{ref}$$

Backup: ML Model Input Features

Aircraft State

Normalized speed deviation from cruise speed; radial position from sector center; route distance; variance of relative bearing and speed among neighbors; angle between planned heading and mean neighbor heading

MVP Resolution

Suggested heading change $\Delta\psi$; suggested speed change magnitude $\| \Delta v \|$; sector congestion N/(πR^{2}sector); neighborhood conflict density.

Conflict Severity

Weighted urgency (w_{deg}) combining temporal and spatial proximity; minimum tCPA; minimum dCPA.

$$w_{\mathrm{deg}} = \frac{1}{N-1} \sum_{i \in \mathcal{C}} \exp\left(-\frac{t_{\mathrm{CPA},i}}{0.35 t_{\mathrm{look}}}\right) \cdot max\left(\frac{r_{\mathrm{pz}} - d_{\mathrm{CPA},i}}{r_{\mathrm{pz}}}, 0\right)$$

Backup: Training Hyperparameters

Component	Value
Hidden dimension	128
FFN blocks	4
Activation	SiLU
Dropout	0.05
Log-variance clamp	$[-8, 3]$
Optimizer	AdamW
Learning rate	3×10^{-4}
Weight decay	10^{-4}
Batch size	256
Gradient clipping	1.0
Epochs	10,000
Train/validation split	80/20
Selection criterion	Min validation NLL

Backup: Energy Overhead Prediction Statistics Per Density

N	Mean (%)	Median (%)	P90 (%)	P95 (%)	Max (%)
10	1.42	0.98	3.36	4.51	6.98
20	1.34	0.94	3.06	3.84	12.56
30	1.43	1.00	3.07	3.87	24.87
40	1.61	1.13	3.52	4.38	22.97
50	1.85	1.29	3.76	4.72	36.49
60	2.12	1.44	4.06	5.30	44.23

Key observations:

- Median remains remarkably stable (<1.5%) across all densities
- P95 grows modestly from 4.5% to 5.3%
- Maximum overhead increases substantially (7% → 44%) due to cascading conflicts

Backup: Why Not Deep RL for the Conflict Resolution

Aspect	MVP	Deep RL
Explainability	✓ Geometric, deterministic	✗ Black-box
Certification	✓ Auditable decision logic	✗ Challenging
Computational	✓ Lightweight	✗ Inference cost
High-density perf.	✓ Validated	✓ Promising
Energy characterization	✓ This paper	? Unknown

MVP provides a well-understood baseline for energy characterization. Our framework could be applied to characterize RL-based methods in future work.