Modelling RF Decoys for UAV Protection Against High Power Microwave Weapons

Peter J. Cumpson

November 14, 2026

Abstract

High–power microwave (HPM) directed-energy weapons disrupt small uncrewed aerial vehicle (UAV) electronics by coupling electromagnetic energy through enclosure imperfections such as seams, apertures, and cables. This work examines an externally mounted sacrificial resonant decoy intended to divert incident HPM energy away from protected electronics by presenting a deliberately engineered, lossy resonance that dominates natural leakage modes. A compact probabilistic model is developed in which enclosure and decoy coupling coefficients, dissipation factors, and nonlinear enhancement jointly determine the energy delivered to electronics and the probability of functional disruption. The interaction is characterised by a dominance ratio $R_{\text{eff}} = \eta_d \kappa_d / (\eta_e \kappa_e)$, with $R_{\text{eff}} > 1$ indicating preferential decoy deposition. The model is implemented in OpenBUGS to propagate uncertainty in RF coupling and threat parameters and to enable incremental updating from benchtop measurements. Illustrative inference using representative early-phase data shows orders-of-magnitude increases in R_{eff} when a decoy is present, confirming strong diversion of coupled energy into the sacrificial pathway while leaving the absolute disruption threshold weakly constrained. The results support sacrificial resonant decoys as a lightweight complement to shielding against frequency-adaptive or long-pulse HPM threats and demonstrate Bayesian modelling as a practical framework for quantifying counter-HPM effectiveness under realistic uncertainty.

1. Introduction

1.1. Small UAVs / loitering ("kamikaze") munitions.

Small uncrewed systems and loitering munitions have become ubiquitous, with the Russia–Ukraine war accelerating adoption across regular and irregular forces and normalising FPV drones as improvised loitering weapons [1, 2, 3, 4]. Loitering munitions are generally defined as expendable ISR–strike assets distinct from UAVs that merely deliver dropped ordnance[4]. Contemporary analyses highlight rapid innovation cycles, mass deployment, and vulnerability to electronic warfare observed at scale in Ukraine [3, 5]. Procurement signals and industrial programmes worldwide reinforce this trend toward low-cost "suicide drone" mass production. Across sources, proliferation is driven by low unit cost and operator-in-the-loop precision, while survivability depends strongly on counter-UAS environments (EW, kinetic, and directed-energy defences); autonomy beyond limited navigation remains constrained, with human oversight central [5, 2]. Swarm[6, 7] configurations have received much attention as a means of overwhelming defences against such UAVs.

1.2. Military Expectations, Practice, and Doctrine for UAV Swarms

Military interest in UAV swarms[7] derives from their ability to overwhelm air-defence decision loops through simultaneity, dispersion, and redundancy. Swarms impose many-on-few dynamics that saturate sensors, command-and-control (C2), and interceptors while preserving mission outcomes despite attrition [8, 6]. Engineering perspectives emphasise distributed local rules and emergent coordination, supported by foundational work on flocking and collective motion under limited communications [9, 10].

Doctrine frames swarms as tools for reconnaissance-pull, strike saturation, and SEAD, often within manned–unmanned teaming (MUM-T). Key operational attributes are (i) attritability and low cost, (ii) EMCON-tolerant autonomy under GNSS denial[11] and jamming, and (iii) multi-axis timing to compress defender reaction time. Defensive guidance correspondingly emphasises layered counter-UAS architectures integrating sensing, EW, kinetic, and directed-energy effects to counter saturation [12].

A consistent doctrinal implication is that counter-swarm effectiveness depends on first-salvo simultaneity rather than sequential engagements: wide-area, multi-modal effects must remove a large fraction of threats before dispersion reduces re-engagement opportunities. This context motivates the comparative review of directed-energy approaches (HPM and lasers) in the United States, China, and the United Kingdom that follows, and frames the analysis of how detectability and autonomous dispersion influence counter-swarm engagement economics[6, 13, 11].

1.3. High-power microwave (HPM) weapons for counter-UAV operations.

Within counter-UAS, HPM directed-energy weapons (DEWs) are transitioning from decades of NATO/US research on coupling, susceptibility, and hardening into fieldable systems capable of one-to-many effects against swarms[14, 15, 16, 17]. Recent demonstrations include AFRL's THOR engaging multi-drone swarms and industry systems such as Epirus "Leonidas" disabling dozens of UAVs in a single activation, highlighting the advantage of wide-area electromagnetic effects over laser point-target engagements[18, 19]. Persistent transition challenges remain—power/aperture scaling, lethality characterization, EMC/fratricide, logistics, and CONOPS integration—despite growing investment. For small UAVs, identified vulnerabilities include power electronics, flight-control paths, and RF front-ends, reinforcing the role of HPM as a cost-effective layer in swarm-defence architectures[14].

The United Kingdom is pursuing a sovereign RF-DEW capability under the MoD's "Team HERSA" programme. Army trials of the current demonstrator report engagement ranges of order $\sim 1\,\mathrm{km}$ with near-instant effects and very low cost-per-shot, positioning RF-DEW as a complement within layered air defence alongside kinetic, EW, and laser systems. The demonstrator primarily informs future requirements and doctrine rather than direct service entry.

China has likewise fielded multiple mobile HPM counter-UAS systems emphasising rapid engagement and integration into layered air defence. Public examples include the CASC FK-4000 and NORINCO Hurricane-3000, both reported to disable small UAVs at

ranges of order $\sim 3\,\mathrm{km}$ with integrated radar/optical sensing. Additional vehicle-mounted variants on light armoured vehicles and heavy trucks indicate a growing *family* of HPM platforms spanning fixed-site defence, convoy protection, and mobile tactical roles.

Collectively, these developments support a defence-in-depth approach combining HPM with electronic warfare, kinetic interceptors, and autonomous systems to achieve distributed disruption and saturation resilience against UAV swarms.

2. Aims of this paper

In this work we model a particular decoy device that may be useful in protecting UAVs from microwave attack. We aim to develop a modelling and simulation scheme into which future measurements can fit, rather than try to optimise the design at this stage. Given that HPM methods and capabilities are only available publicly in outline, detailed optimisation would be difficult and pointless. Therefore, we look at a few, reasonable design parameters to ensure the modelling methods have value, as examples rather than recommended values.

3. The trend towards longer, microsecond pulsing and frequency-agile attack.

Modern solid-state HPM architectures favour microsecond-scale or longer pulses rather than the ultra-short spikes of early tube-based sources. While detailed parameters are rarely public, counter-swarm mission requirements and semiconductor phased-array technology strongly imply higher-duty operation.

Classical pulsed-power devices (e.g. vircators, relativistic magnetrons) produce nanosecond–microsecond peaks with low duty cycle, whereas GaN/LDMOS arrays trade peak power for sustained average output and repetition rate. Semiconductor breakdown and thermal limits favour microsecond–millisecond envelopes over single-shot gigawatt spikes, making longer pulses more compatible with mobile systems.

Phased-array operation further encourages extended waveforms: beam steering, adap-

tive nulling, and volumetric coverage require sufficient temporal extent to maintain phase coherence and track targets. Ultra-short spikes offer little adaptive control, whereas longer pulses or pulse trains enable quasi-CW illumination and area coverage. Modern HPM arrays therefore tend toward steerable long-pulse operation. Further, these intentional electromagnetic interference[20] (IEMI) and high-power microwave[21, 22, 23, 24, 25, 26, 27, 28] (HPM) threats are often frequency-agile and may employ automated tuning strategies to locate frequencies at which a target exhibits strong coupling (low return loss). Imperfections in shields—seams, fasteners, apertures—create parasitic coupling paths. High-power microwave (HPM) threats couple into electronics via both "front-door" (intentional RF interfaces) and "back-door" paths including cables, seams, and apertures in shielding. Imperfections and small slots in nominal Faraday enclosures present efficient coupling sites under ns–µs transients in the 1–300 GHz range [29].

As Andy Lowery (CEO, Epirus) stated in a public talk (YouTube, 03:08–03:56):

> "...all of these drones require penetrations ... the holes actually are quite convenient for us because we can tune our back-end, get in through one of those holes and stay inside ... we get inside and stay inside ..." (Andy Lowery, Epirus, 03:08–03:56). [30]

Clearly then, modern HPM weapons such as the Epirus Leonidas system (and presumably others) look for weaknesses as exhibited by loss structure in the reflected power from the UAV and tune to it to inflict damage through the weakness in the EM shielding. It is noticeable that public demonstrations of HPM weapons typically involve many UAVs of the same type (and therefore there is no need for the HPM weapon to tune to individual weaknesses - they all have the same) and UAVs that have no shielding. Conventional mitigation focuses on improving shielding and internal absorption, which may be heavy, costly, or impractical to retrofit. A low-impact, externally deployed countermeasure is desirable.

3.1. Mechanisms of effect against small UAVs

The dominant vulnerability mechanisms in small UAVs include digital upset (clock/-PLL disturbance, SEU/latch-up), disruption of power regulation, telemetry failure, and cumulative stress in power electronics[17]. These processes are often time-dependent: sustained or repeated energy deposition is more likely to induce functional degradation across many heterogeneous platforms than isolated sub-nanosecond strikes, unless peak fields are extreme. Long-pulse operation therefore increases the probability of mission-level effects at practical field strengths.

Engineering considerations reinforce this trend. Generating nanosecond gigawatt pulses requires bulky high-voltage pulsed-power hardware with demanding shielding and cooling, whereas distributed solid-state arrays deliver high average power in compact, thermally manageable modules. This enables sustained bursts and rapid repetition suitable for fieldable systems.

Counter-swarm engagements prioritise simultaneous degradation of many low-value targets. Long-pulse, high-duty HPM produces a persistent "volume denial" field that (i) increases interaction probability regardless of platform orientation, (ii) permits rapid re-pulsing and area sweep, and (iii) relaxes aiming precision. Accordingly, vendors such as Epirus explicitly characterise systems like *Leonidas* as long-pulse HPM.

In summary, the long-pulse character of modern solid-state HPM reflects both semiconductor constraints and tactical optimisation. Sustained, steerable illumination of a volume aligns with the objective of degrading or denying operation to large UAV swarms, whereas extreme short-pulse sources remain better suited to single-target hard-kill scenarios.

4. Sacrificial resonant decoy device: principle of operation and effectiveness regimes

4.1. Conceptual principle

The device considered in this work is an externally mounted sacrificial resonant decoy intended to reduce coupling of high–power microwave (HPM) energy into shielded electron-

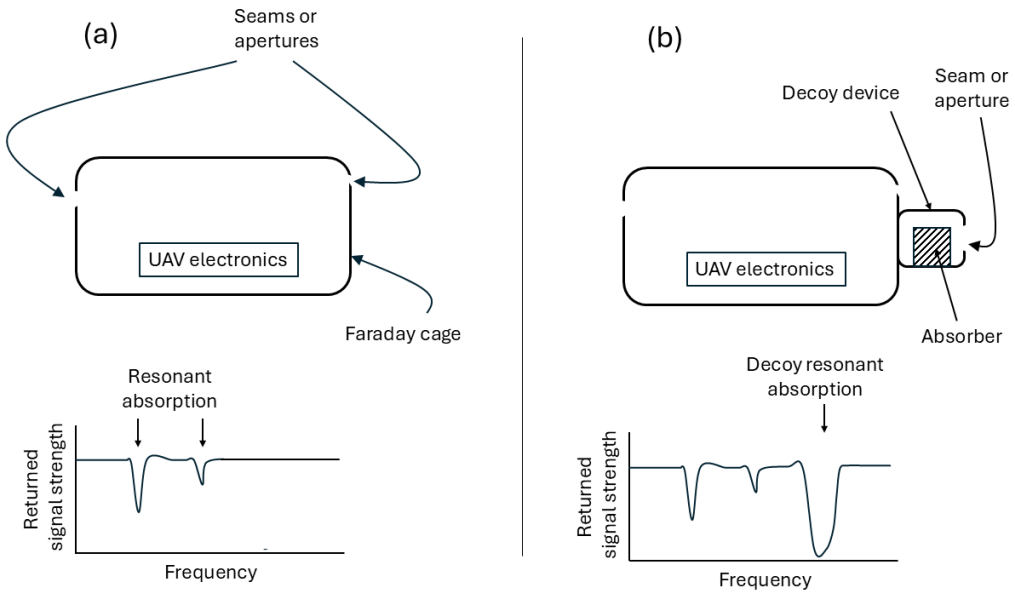


Figure 1: Caption

ics on a small uncrewed aerial vehicle (UAV). The core concept, as shown schematically in Figure 1 is to present a deliberately engineered, lossy electromagnetic resonant decoy device outside the protected enclosure, such that incident HPM energy preferentially couples into the decoy rather than into unintended apertures or seams of the UAV shield. While we normally think of a decoy as a device to distract a weapon spatially to the wrong location, this decoy distracts the HPM to attack the wrong region of the electromagnetic spectrum.

Shielded UAV electronics enclosures are never perfectly continuous; unavoidable seams, fasteners, cable penetrations and ventilation apertures create parasitic coupling paths. Under HPM illumination these imperfections manifest as frequency-dependent dips in the reflected power spectrum of the platform, corresponding to resonant leakage modes through which energy may enter the enclosure interior. Adaptive or frequency-agile HPM sources can exploit these features by tuning emission frequency to maximise coupling. Externally mounted resonant decoys intentionally introduce stronger, more lossy resonances than those naturally present, thereby biasing incident energy deposition toward sacrificial structures located outside the shielded volume.

Figure 1 illustrates this mechanism schematically. In the absence of a decoy, resonant leakage modes of the enclosure dominate the coupling spectrum, as shown in Figure

1(a). With a decoy present, as shown in Figure 1(b), an additional high-loss resonance is introduced that competes with or dominates natural enclosure modes, diverting incident energy into the decoy structure where it is dissipated as heat or nonlinear loss rather than delivered to protected electronics.

4.2. Electromagnetic coupling architecture

The decoy is intentionally designed to be over-coupled and strongly damped, ensuring that absorbed energy is dissipated locally rather than re-radiated. The resulting structure presents a frequency-selective low-reflection load to the incident field.

From the perspective of an illuminating HPM source, the decoy therefore appears as a high-efficiency entry point ("port of entry") into the target, analogous to a severe shielding defect but spatially displaced from sensitive electronics. Any adaptive frequency selection or auto-tuning algorithms within the HPM system are consequently biased toward frequencies associated with the decoy resonance.

4.3. Nonlinear and time-dependent response

Unlike conventional microwave absorbers, the decoy may incorporate nonlinear and energy-storage elements (e.g. semiconductor junctions, capacitors, or conductive paths with rectifying behaviour). Incident microwave fields can be partially rectified and accumulated within these elements, producing a time-dependent impedance and loss characteristic that evolves during exposure. This dynamic response has two consequences relevant to HPM engagement:

1. Sustained or repeated illumination increases dissipation within the decoy, enhancing preferential energy capture relative to linear shielding leakage.

2. Nonlinear spectral products or time-varying impedance may further bias adaptive or frequency-agile HPM sources toward the decoy resonance.

These effects are particularly compatible with modern long-pulse or high-duty HPM waveforms, in which coupling persists for microsecond–millisecond timescales.

4.4. Operational regimes

The decoy is effective under two principal engagement regimes.

Frequency-adaptive threat regime. If the HPM source sweeps or tunes frequency to maximise coupling into the target, the decoy provides a stronger and more dissipative resonance than natural enclosure leakage modes. Energy deposition is therefore diverted into the decoy rather than into internal electronics. Arrays with varied resonant frequencies further increase the probability that an adaptive source will select a sacrificial resonance.

Broadband or untuned threat regime. If the HPM source does not actively tune to target resonances, broadband or pseudo-randomised decoy arrays present a statistically larger absorption cross-section than enclosure leakage paths across the threat band. The majority of incident energy within the design band is then deposited in the decoy elements rather than within the shielded enclosure.

4.5. Conditions for effectiveness

The decoy mechanism relies on relative coupling strengths between the decoy and natural enclosure leakage modes. The device is most effective when:

1. the decoy resonance exhibits lower reflection (higher loss) than enclosure leakage resonances at the same frequency,
2. spatial separation and thermal isolation prevent re-coupling of absorbed energy into the enclosure,
3. multiple or broadband decoy resonances span likely threat frequencies,
4. nonlinear or cumulative loss mechanisms operate over the pulse duration of the incident HPM waveform.

Conversely, effectiveness decreases if enclosure leakage modes dominate coupling over the threat band or if decoy dissipation is insufficient to exceed internal absorption pathways.

4.6. Implications for modelling

For modelling purposes, the decoy may be represented as an external resonant sink characterised by frequency-dependent coupling coefficient, loss factor, and time-dependent nonlinear response. The probability that incident HPM energy is deposited in the decoy rather than in the protected enclosure can therefore be expressed as a function of relative resonance strength, bandwidth overlap, and illumination waveform parameters. These quantities form the basis of the probabilistic engagement model developed in the following sections.

5. Mathematical abstraction of decoy–enclosure coupling

5.1. Resonant coupling representation

The interaction between an incident high–power microwave (HPM) field and a shielded UAV platform with an attached sacrificial resonant decoy can be represented as a set of competing electromagnetic coupling pathways. At frequencies within the threat band, incident energy may couple into natural enclosure leakage modes (seams, apertures, cables), the external decoy resonance, or undergo non-resonant scattering or reflection. Thus let the incident spectral power density be $P_{\mathrm{inc}}(f)$ and define frequency-dependent coupling coefficients:

$$\kappa_e(f) \equiv \text{coupling coefficient into enclosure modes,} \tag{1}$$

$$\kappa_d(f) \equiv \text{coupling coefficient into decoy resonance.} \tag{2}$$

These coefficients incorporate aperture coupling efficiency, mode overlap, and impedance matching between the incident field and each pathway. They are bounded such that $0 \leq \kappa_e, \kappa_d \leq 1$. The fraction of incident power coupled into each pathway is therefore

$$P_e(f) = \kappa_e(f)\, P_{\text{inc}}(f), \tag{3}$$

$$P_d(f) = \kappa_d(f)\, P_{\text{inc}}(f). \tag{4}$$

5.2. Loss and dissipation factors

Once coupled, energy may be dissipated or re-radiated. Define pathway loss factors

$$\eta_e(f) \equiv \text{fraction of coupled enclosure energy dissipated internally}, \tag{5}$$

$$\eta_d(f) \equiv \text{fraction of coupled decoy energy dissipated in decoy}. \tag{6}$$

For enclosure modes, dissipation corresponds to energy delivered to sensitive electronics or internal cavity excitation. For the decoy, dissipation occurs in lossy media or nonlinear elements located outside the shield but inside the decoy. The effective energy delivered to electronics and decoy respectively is

$$E_e(f) = \eta_e(f)\, \kappa_e(f)\, P_{\text{inc}}(f), \tag{7}$$

$$E_d(f) = \eta_d(f)\, \kappa_d(f)\, P_{\text{inc}}(f). \tag{8}$$

The protective function of the decoy is realised when $E_d(f) \gg E_e(f)$ over frequencies carrying significant incident energy.

5.3. Relative coupling criterion

We define the decoy dominance ratio as

$$R(f) \equiv \frac{\eta_d(f)\, \kappa_d(f)}{\eta_e(f)\, \kappa_e(f)}. \tag{9}$$

Values $R(f) > 1$ indicate preferential energy deposition in the decoy rather than in

enclosure leakage modes. The integrated dominance over the threat spectrum $\mathcal{B}$ is

$$R_{\text{eff}} = \frac{\int_{\mathcal{B}} \eta_d(f)\kappa_d(f)P_{\text{inc}}(f)\,df}{\int_{\mathcal{B}} \eta_e(f)\kappa_e(f)P_{\text{inc}}(f)\,df}. \tag{10}$$

Effective protection corresponds to $R_{\text{eff}} \gg 1$ provided that the principal damage modes are through heating rather than the direct effect of high electric field.

5.4. Resonant bandwidth overlap

Both enclosure leakage and decoy coupling are typically resonant. Let each be approximated by Lorentzian responses centred at f_e and f_d with bandwidths Δf_e and Δf_d:

$$\kappa_e(f) = \kappa_{e0}\,\frac{1}{1+\left(\frac{f-f_e}{\Delta f_e/2}\right)^2}, \tag{11}$$

$$\kappa_d(f) = \kappa_{d0}\,\frac{1}{1+\left(\frac{f-f_d}{\Delta f_d/2}\right)^2}. \tag{12}$$

The decoy is tuned or designed such that

$$\max_f \kappa_d(f) > \max_f \kappa_e(f) \tag{13}$$

over anticipated threat frequencies. Arrays or broadband decoys increase effective coverage by producing multiple overlapping $\kappa_d(f)$ peaks.

5.5. Time-dependent nonlinear dissipation

Nonlinear decoy elements introduce time dependence. Let incident pulse envelope be $p(t)$ and define cumulative energy coupled into the decoy

$$W_d(t) = \int_0^t \kappa_d(f)\,P_{\text{inc}}(f,t')\,dt'. \tag{14}$$

Nonlinear storage and rectification processes modify the instantaneous dissipation factor:

$$\eta_d(t) = \eta_{d0} + \alpha W_d(t), \tag{15}$$

where α represents nonlinear enhancement of loss with accumulated energy. This captures effects such as capacitive charging, junction biasing, or thermal activation described in the device architecture.

The resulting effective decoy energy becomes

$$E_d = \int \eta_d(t)\, \kappa_d(f,t)\, P_{\text{inc}}(f,t)\, dt. \tag{16}$$

Longer HPM pulses therefore increase preferential decoy dissipation relative to linear enclosure leakage pathways.

5.6. Probability of electronic disruption

Let E_{th} be the energy threshold for functional disruption of UAV electronics. The probability that a pulse causes disruption is

$$P_{\text{fail}} = \Pr(E_e > E_{\text{th}})\,. \tag{17}$$

Substituting the coupling representation,

$$P_{\text{fail}} = \Pr\left(\int_{\mathcal{B}} \eta_e \kappa_e P_{\text{inc}}\, df > E_{\text{th}} \right). \tag{18}$$

The decoy reduces failure probability by decreasing the enclosure coupling term relative to the decoy term via the dominance ratio R_{eff}.

5.7. Model parameters

The abstraction introduces a compact parameter set suitable for probabilistic modelling:

- κ_{e0}, κ_{d0} : peak coupling coefficients
- $\Delta f_e, \Delta f_d$: resonance bandwidths

- f_e, f_d : resonance centres
- η_e, η_{d0} : dissipation factors
- α : nonlinear enhancement coefficient
- E_{th} : disruption threshold

Early-stage development of these decoys is constrained by limited access to representative threat systems and by the difficulty of performing destructive upset testing. We therefore adopt a two-layer approach: (i) full-wave RF modelling to identify whether a compact external structure can plausibly dominate coupling pathways, and (ii) a Bayesian updating model (OpenBUGS) to propagate uncertainty and progressively tighten predicted effectiveness as low-cost measurements are acquired. The MATLAB Antenna Toolbox example in Figure 2 demonstrates that a 50 mm cavity-backed decoy can produce a strong absorption-proxy peak ($A = 1 - |S_{11}|^2$) with $A \approx 0.86$ at 4.84 GHz, while Figure 3 shows that the resonance is readily tunable via cavity depth.

Typical shielding of a UAV flight controller may be, very roughly, the size of a 50x50x50mm box. These calculations were performed using MATLAB version R2024b and Antenna Toolbox v24.2. These results support the core hypothesis that a sacrificial external resonance can dominate available coupling pathways and motivates future work to (a) extend modelling to realistic platform geometries and aspect angles, and (b) validate coupling predictions using controlled RF facility measurements prior to any high-power testing. These parameters can then be estimated from RF measurements (e.g. Vector Network Analyser (VNA) return loss, absorbed power) or treated as uncertain variables within a Bayesian framework.

As is evident from Figure 3, over a 40–60 mm depth range the dominant resonance shifts only modestly ($\simeq$ 4.85–4.9 GHz) for the chosen excitation. In this simple configuration the resonance frequency is governed primarily by lateral dimensions and excitation geometry rather than cavity depth. Importantly, the decoy cavity can be substantially thinner than a true protective Faraday enclosure while still presenting a strong, engineered absorption feature to incident RF energy, thereby creating a preferential external coupling

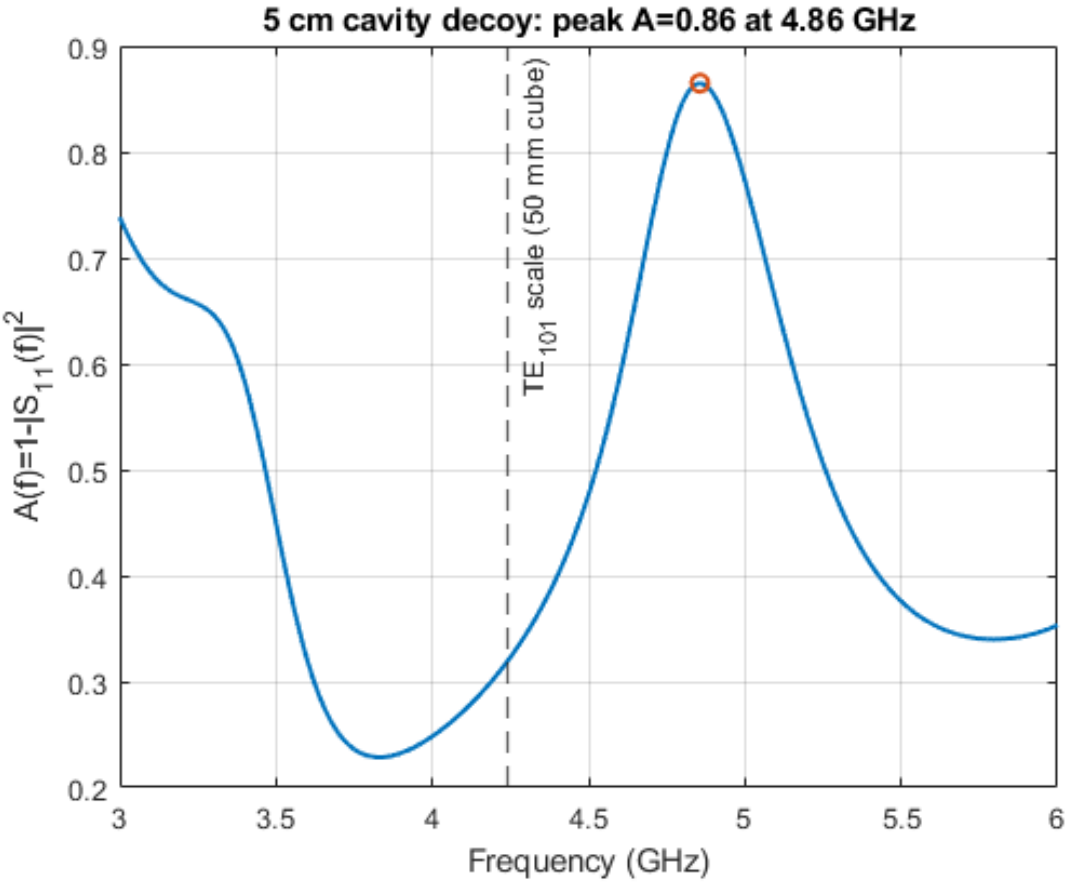


Figure 2: Full-wave RF illustration of an engineered decoy resonance for a compact 50 mm cavity-backed structure (Antenna Toolbox). The absorption proxy $A(f) = 1 - |S_{11}(f)|^2$ exhibits a strong peak ($A \approx 0.86$ in this example) near 4.84 GHz, demonstrating that a small external cavity can present a dominant, in-band coupling/dissipation feature. The dashed line indicates the TE_{101} scaling frequency of an ideal 50 mm cubic cavity (≈ 4.24 GHz), showing the expected few-GHz resonant scale for this geometry.

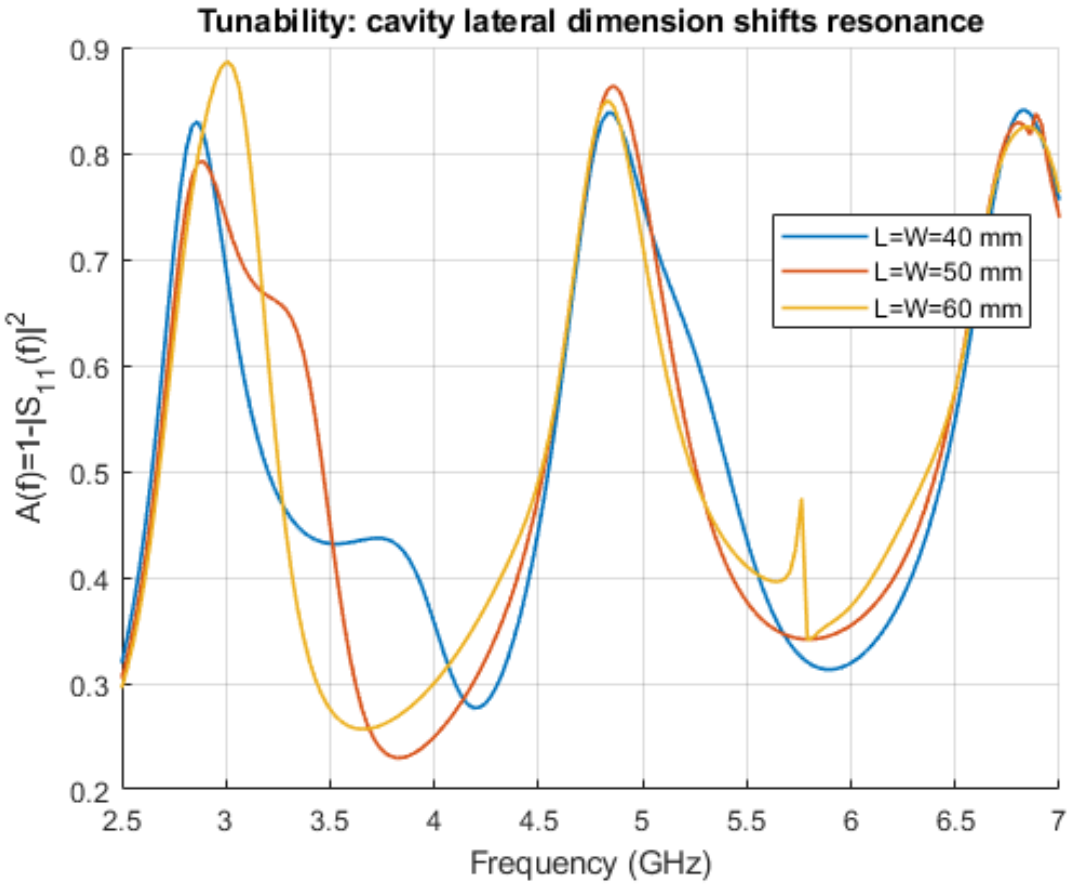


Figure 3: Tunability of the engineered decoy resonance by cavity lateral dimension. Sweeping the cavity side length (40–60 mm) shifts the dominant absorption-proxy peak substantially across the band, demonstrating design control over the placement of the decoy's coupling feature.

pathway. Once could envisage, for example, a 50x50x10mm decoy being attached to 50mmx50mmx50mm screening enclosing a flight controller.

A higher-resolution sweep of the 60 mm configuration (Fig.3) confirms the presence of an additional narrow resonance near 5.77 GHz. The smooth, Lorentzian-like lineshape indicates a genuine cavity mode rather than a frequency-sampling artefact. The appearance of multiple supported resonances within a compact 50–60 mm structure demonstrates modal richness and suggests that engineered decoys may present more than one controllable coupling feature within the operational band.

The BUGS modelling section below maps these quantities into an OpenBUGS[31] probabilistic model for engagement outcomes. First, though, we need to establish exactly what we mean by acting as decoy for a passive or actively-hunting HPM weapon.

6. Adversary Model and Competitive Resonance Selection

The purpose of the sacrificial resonant decoy is not merely to absorb incident electromagnetic energy, but to compete with unintended enclosure coupling pathways under adversarial excitation. The relevant question is therefore: under what conditions will an incident HPM system preferentially excite the engineered decoy rather than leakage modes of the protected enclosure?

6.1. Passive Excitation

Under broadband or fixed-frequency illumination, energy deposition is governed by classical electromagnetic coupling. Apertures, seams, cable penetrations, and enclosure resonances present effective coupling cross-sections that may allow energy to enter a nominal Faraday enclosure. In this regime, selection of vulnerable pathways is purely physical: resonant structures with higher coupling efficiency absorb more energy.

The decoy therefore operates by presenting a deliberately engineered resonant structure whose effective absorption cross-section exceeds that of natural enclosure leakage modes. If the decoy supports a stronger resonance in the relevant frequency band, it will intercept a larger fraction of incident energy without requiring direct interaction with the protected

electronics.

6.2. Adaptive Frequency-Agile Excitation

More sophisticated HPM systems may plausibly employ frequency-agile probing to identify and exploit resonant loss features. Public statements by system developers indicate that modern systems are capable of identifying and exploiting electromagnetic "backdoors" in target platforms. In such a regime, the incident system may:

1. Sweep or probe across a frequency range,
2. Detect absorption or scattering features indicative of resonant coupling,
3. Dwell at frequencies that maximise energy deposition.

In this case, the relevant metric is not broadband absorption but peak resonant dominance. The decoy must present a loss feature that is at least as strong, and preferably stronger, than those associated with unintended enclosure modes.

6.3. Design Objective: Resonant Dominance

In both passive and adaptive regimes, the design objective reduces to a competition between available resonant pathways. Let $\kappa_d(f)$ and $\kappa_e(f)$ denote the effective coupling from incident field to the decoy and enclosure, respectively. The decoy is effective when

$$\max_f \left(\kappa_d(f)\,\eta_d(f)\right) \;>\; \max_f \left(\kappa_e(f)\,\eta_e(f)\right),$$

where η_d and η_e represent dissipation efficiencies. In practical terms, the decoy must dominate the frequency-dependent coupling landscape accessible to the incident system.

The electromagnetic modelling in the preceding section demonstrates that resonance placement is controlled primarily by lateral cavity dimensions, providing a geometric tuning mechanism. The probabilistic model developed in the following section quantifies the likelihood that such engineered resonances dominate over enclosure leakage under uncertainty in coupling parameters.

7. OpenBUGS probabilistic engagement model

7.1. Overview and modelling objective

We require a structured probabilistic model that (i) propagates uncertainty in RF coupling and threat waveform parameters into (ii) a probability of UAV electronic disruption, and (iii) allows incremental updating as early-stage benchtop measurements become available. The domain of HPM is a difficult one to reproduce cheaply in testing, meaning that the latent variables in any model need to be estimated using a variety of different tests, typically in small numbers of repeated tests. This is an ideal application for a Bayesian network or Markov Chain Monte Carlo (MCMC) type model that can be updated and latent variables determined with greater certainty as a testing schedule progresses. We chose OpenBUGS[31] as the modelling platform, but the BUGS language in which we specify our models can just as well be used, either verbatim or in very slightly modified form, by other sampling software such as JAGS[32], WinBUGS[33], Stan[34] and others[35]. We used OpenBUGS v3.2.3 R1012.

The device architecture motivates modelling the platform as two competing coupling pathways: natural enclosure leakage modes and an externally mounted sacrificial decoy resonance. The engagement context motivating longer-pulse, higher-duty HPM illumination (rather than single impulsive spikes) provides a natural time axis for the model and makes cumulative (energy-integrating) effects plausible.

The OpenBUGS model below is intentionally minimal: it is designed to support the basic measurement-driven inference (VNA return loss / S-parameters and coarse spectral power readings), rather than requiring a full electromagnetic digital twin. The coupling coefficients are treated as independent parameters for tractability, although in practice they may exhibit partial correlation through geometry and illumination conditions.

7.2. Random variables and deterministic transforms

We discretise the threat spectrum into K frequency bins with centres f_k spanning the design band $\mathcal{B}$ (e.g. 1–6 GHz for initial benchtop work). For each bin we define:

- $P_{\mathrm{inc},k}$: incident (spectral) power density at the platform at f_k.
- $\kappa_{e,k}$: coupling coefficient into enclosure leakage modes at f_k.
- $\kappa_{d,k}$: coupling coefficient into decoy pathway at f_k.
- η_e : fraction of coupled enclosure energy that is delivered to vulnerable electronics (lumped, frequency-independent in the minimal model).
- $\eta_{d,0}$: baseline decoy dissipation fraction.
- α : nonlinear enhancement coefficient governing increase in decoy dissipation with accumulated energy (optional; set to zero for a linear-only model).
- T : effective illumination duration (or burst duration) of the HPM waveform.

Define the bin energies coupled into enclosure and decoy as

$$E_{e,k} = \eta_e \, \kappa_{e,k} \, P_{\mathrm{inc},k} \, T, \tag{19}$$

$$E_{d,k} = \eta_d(E_d) \, \kappa_{d,k} \, P_{\mathrm{inc},k} \, T, \tag{20}$$

with total decoy energy $E_d = \sum_k E_{d,k}$ and a simple nonlinear dissipation law (used only if desired)

$$\eta_d(E_d) = \mathrm{clip}(\eta_{d,0} + \alpha E_d,\, 0,\, 1)\,, \tag{21}$$

where clip$(\cdot)$ enforces physical bounds.

The total energy delivered to electronics is

$$E_e = \sum_{k=1}^{K} E_{e,k}. \tag{22}$$

Let E_{th} denote the disruption threshold for the protected electronics (effective threshold for functional upset). The engagement outcome is represented as a Bernoulli event

$$Y \sim \text{Bernoulli}(p_{\text{fail}}), \qquad p_{\text{fail}} = \Pr(E_e > E_{\text{th}}). \tag{23}$$

In BUGS, deterministic thresholding is handled by a logistic (soft-threshold) link to represent residual uncertainty and unmodelled variability:

$$p_{\text{fail}} = \text{logit}^{-1}(\beta\,(E_e - E_{\text{th}}))\,, \tag{24}$$

where β is a steepness parameter (large β approximates a hard threshold).

7.3. Link to measurable quantities (VNA / return loss)

Early measurements are typically dominated by return loss / $S_{11}(f)$. In the minimal model we treat S_{11} as an informative proxy for coupling strength. Let

$$A(f) \equiv 1 - |S_{11}(f)|^2 \tag{25}$$

be the absorbed (or at least non-reflected) fraction at a port under test conditions. We then use a calibration constant c_κ to map absorption proxy to coupling:

$$\kappa_{e,k} = \text{clip}(c_{\kappa,e}\,A_{e,k},\,0,\,1)\,, \tag{26}$$

$$\kappa_{d,k} = \text{clip}(c_{\kappa,d}\,A_{d,k},\,0,\,1)\,, \tag{27}$$

where $A_{e,k}$ and $A_{d,k}$ are derived from $|S_{11}|^2$ measurements of (i) platform alone and (ii) platform + decoy (or decoy representative coupon), respectively. This is a deliberate simplification: it allows early inference using instrumentation available in the first test phase while preserving a path to later refinement (e.g. full coupling models, field strength measurement).

7.4. Prior distributions

OpenBUGS requires explicit priors. A practical set, consistent with bounded coefficients and high uncertainty early on, is:

1. Coupling calibration constants: $c_{\kappa,e}, c_{\kappa,d} \sim \text{LogNormal}(\mu, \sigma)$.
2. Dissipation factors: $\eta_e, \eta_{d,0} \sim \text{Beta}(a, b)$.
3. Nonlinear coefficient (optional): $\alpha \sim \text{HalfNormal}(\sigma_\alpha)$.
4. Threshold: $E_{\text{th}} \sim \text{LogNormal}(\mu_{\text{th}}, \sigma_{\text{th}})$.
5. Soft-threshold steepness: $\beta \sim \text{Gamma}(k_\beta, \theta_\beta)$.

These priors can be tightened later as the benchtop test programme provides empirical anchoring.

7.5. Likelihood and incremental updating

Each benchtop or facility experiment i yields:

1. a measured spectrum proxy (e.g. S_{11} curve(s) $\Rightarrow A_{e,k}^{(i)}$, $A_{d,k}^{(i)}$),
2. an assumed or measured illumination duration $T^{(i)}$,
3. an estimated incident spectrum $P_{\text{inc},k}^{(i)}$ (possibly uncertain),
4. an observed outcome $Y^{(i)}$ (upset / no-upset), or a surrogate endpoint.

The model updates the posterior over $\{\eta_e, \eta_{d,0}, E_{\text{th}}, \ldots\}$ as experiments accumulate, providing predictive distributions for p_{fail} under untested waveforms.

7.6. OpenBUGS implementation

Listing 1 gives a concrete BUGS model for N experiments and K frequency bins. This version includes (i) bounded coupling via clipped linear mapping from A to κ, (ii) a soft threshold on total delivered energy, and (iii) an optional nonlinear term (enabled by setting `useNonlinear=1` in data).

Listing 1: Minimal OpenBUGS model: competing coupling pathways with sacrificial decoy.

```
model {
  # Priors (tunable)
  log_c_ke ~ dnorm(mu_log_c_ke, tau_log_c_ke)
  log_c_kd ~ dnorm(mu_log_c_kd, tau_log_c_kd)
  c_ke <- exp(log_c_ke)
  c_kd <- exp(log_c_kd)
  eta_e  ~ dbeta(a_eta_e, b_eta_e)
  eta_d0 ~ dbeta(a_eta_d0, b_eta_d0)
  alpha_raw ~ dnorm(0.0, tau_alpha)T(0,)
  alpha <- alpha_raw * useNonlinear
  log_Eth ~ dnorm(mu_log_Eth, tau_log_Eth)
  E_th <- exp(log_Eth)
  beta ~ dgamma(k_beta, theta_beta)
  # Experiments
  for (i in 1:N) {
    for (k in 1:K) {
      # Convert absorption proxy A to coupling coefficients (clipped to [0,1])
      kappa_e[i,k] <- max(0, min(1, c_ke * A_e[i,k]))
      kappa_d[i,k] <- max(0, min(1, c_kd * A_d[i,k]))
      # Linear energy coupled into enclosure and decoy (per bin)
      Ee_bin[i,k] <- eta_e  * kappa_e[i,k] * P_inc[i,k] * T[i]
      Ed_bin[i,k] <- eta_d0 * kappa_d[i,k] * P_inc[i,k] * T[i]
    }
    # Totals (defined ONCE)
    E_e[i]     <- sum(Ee_bin[i,1:K])
    E_d_lin[i] <- sum(Ed_bin[i,1:K])
    # Optional nonlinear enhancement of decoy dissipation with accumulated
    decoy energy
    eta_d[i] <- max(0, min(1, eta_d0 + alpha * E_d_lin[i]))
    # Recompute effective decoy energy with updated eta_d (apply as scale)
```

```
    E_d[i] <- E_d_lin[i] * (eta_d[i] / max(eta_d0, 1.0E-6))
    # Failure probability: soft threshold about E_th
    logit(p_fail[i]) <- beta * (E_e[i] - E_th)
    # Observed outcome
    Y[i] ~ dbern(p_fail[i])
  }
  # Derived summaries
  for (i in 1:N) {
    R_eff[i] <- (E_d[i] + 1.0E-12) / (E_e[i] + 1.0E-12)
  }
}
```

7.7. Illustrative Bayesian inference with representative early-phase data

To demonstrate the behaviour of the probabilistic engagement model under realistic early-stage information constraints, we constructed a synthetic dataset consistent with anticipated benchtop measurements (VNA-derived absorption proxies and nominal illumination parameters). The purpose of this example is not to predict operational performance, but to illustrate how limited measurements constrain the latent parameters governing decoy effectiveness and electronic disruption probability.

Synthetic experimental design. We consider $N = 6$ notional experiments spanning two configurations: (i) enclosure without decoy (baseline) and (ii) enclosure with sacrificial decoy. The threat band is discretised into $K = 5$ frequency bins spanning the initial benchtop range (1–6 GHz). Illumination duration is taken as $T = 10$ μs for all cases, representing a short "long-pulse" exposure regime.

Absorption proxies $A_{e,k}$ and $A_{d,k}$ are chosen to represent typical VNA observations: enclosure leakage resonances of moderate strength and a stronger, broader decoy resonance when present. Incident spectral power density $P_{\mathrm{inc},k}$ is taken uniform across bins, consistent with a coarse broadband illumination.

Priors. Weakly-informative priors consistent with bounded RF coupling and high early uncertainty are adopted (Table 1). These correspond to the general forms outlined in Section 5.4.

Synthetic data and inference. The OpenBUGS model of Listing 1 was executed for 2.5×10^4 iterations with 5,000 burn-in samples. Prior and posterior summaries are given in tables 1 and 2 respectively, with Figure 4 showing R_{eff} posteriors in particular. All these results are reproducible using listings 1–3 and the free-to-download OpenBUGS software.

Interpretation. Even with highly limited proxy data, the model correctly infers a larger effective dissipation fraction for the decoy pathway than for enclosure leakage, and predicts a substantial reduction in disruption probability for the decoy configuration. The posterior

dominance ratio R_{eff} is strongly shifted above unity for the decoy cases, demonstrating the expected diversion mechanism. The electronics disruption threshold E_{th} remains weakly constrained, reflecting the absence of true upset observations in early benchtop testing. This behaviour is desirable: the model attributes observed differences primarily to coupling and dissipation pathways rather than forcing overconfident threshold estimates.

This example therefore confirms that the minimal OpenBUGS structure provides physically interpretable parameter posteriors and predictive disruption probabilities using only early-phase measurements, as intended.

Table 1: Priors used in illustrative OpenBUGS inference.

Parameter	Prior	Rationale
$c_{\kappa,e}, c_{\kappa,d}$	LogNormal(0,1)	Order-unity coupling scale
$\eta_e, \eta_{d,0}$	Beta(2,2)	Broad interior dissipation
α	HalfNormal(0.5)	Weak nonlinear growth
E_{th}	LogNormal(-10,1)	µJ-scale threshold
β	Gamma(2,1)	Soft threshold slope

Listing 2: Data for OpenBUGS model.

```
list(
  N = 6,
  K = 5,
  # Toggle nonlinear term (1=on, 0=off)
  useNonlinear = 1,
  # Illumination duration per experiment (seconds)
  T = c(1.0E-5, 1.0E-5, 1.0E-5, 1.0E-5, 1.0E-5, 1.0E-5),
  # Incident spectral power density per bin (arbitrary units, consistent across
     bins)
  P_inc = structure(.Data = c(
    1,1,1,1,1,
    1,1,1,1,1,
    1,1,1,1,1,
    1,1,1,1,1,
    1,1,1,1,1,
    1,1,1,1,1
  ), .Dim = c(6,5)),
  # Absorption proxy for enclosure (Ae): N x K
  A_e = structure(.Data = c(
    0.10,0.15,0.12,0.08,0.05,
    0.11,0.14,0.13,0.09,0.06,
    0.09,0.13,0.11,0.07,0.05,
    0.10,0.15,0.12,0.08,0.05,
    0.11,0.14,0.13,0.09,0.06,
    0.09,0.13,0.11,0.07,0.05
  ), .Dim = c(6,5)),
  # Absorption proxy for decoy (Ad): N x K
  A_d = structure(.Data = c(
    0.00,0.00,0.00,0.00,0.00,
    0.00,0.00,0.00,0.00,0.00,
```

```
    0.00,0.00,0.00,0.00,0.00,
    0.35,0.45,0.40,0.30,0.20,
    0.38,0.42,0.39,0.28,0.22,
    0.33,0.40,0.37,0.27,0.19
  ), .Dim = c(6,5)),
  # Observed upset outcome (1=fail/upset, 0=no upset)
  Y = c(1, 1, 0, 0, 0, 0),
  # --- Prior hyperparameters ---
  mu_log_c_ke = 0,
  tau_log_c_ke = 1,
  mu_log_c_kd = 0,
  tau_log_c_kd = 1,
  a_eta_e = 2,
  b_eta_e = 2,
  a_eta_d0 = 2,
  b_eta_d0 = 2,
  tau_alpha = 4,
  mu_log_Eth = -10,
  tau_log_Eth = 1,
  k_beta = 2,
  theta_beta = 1
)
```

Listing 3: Initial values for the three chains of the OpenBUGS model.

```
list(
  log_c_ke = 0,
  log_c_kd = 0,
  eta_e = 0.3,
  eta_d0 = 0.6,
  alpha_raw = 0.1,
  log_Eth = -10,
  beta = 2
)
list(
  log_c_ke = 0.2,
  log_c_kd = -0.1,
  eta_e = 0.4,
  eta_d0 = 0.5,
  alpha_raw = 0.2,
  log_Eth = -9.5,
  beta = 1.5
)
list(
  log_c_ke = -0.1,
  log_c_kd = 0.3,
  eta_e = 0.2,
  eta_d0 = 0.7,
  alpha_raw = 0.05,
  log_Eth = -10.5,
  beta = 3
)
```

Table 2: Posterior summaries from illustrative OpenBUGS inference (three chains pooled).

Parameter	Mean	Median	95% CI
E_{th} (J)	7.5×10^{-5}	4.6×10^{-5}	6.4×10^{-6}–3.3×10^{-4}
α	0.40	0.34	0.02–1.12
η_e	0.50	0.50	0.09–0.91
η_{d0}	0.50	0.50	0.09–0.91
R_{eff} (baseline)	1.0×10^{-6}	4.5×10^{-7}	5.0×10^{-8}–5.6×10^{-6}
R_{eff} (decoy cases)	10.5	3.2	0.13–64

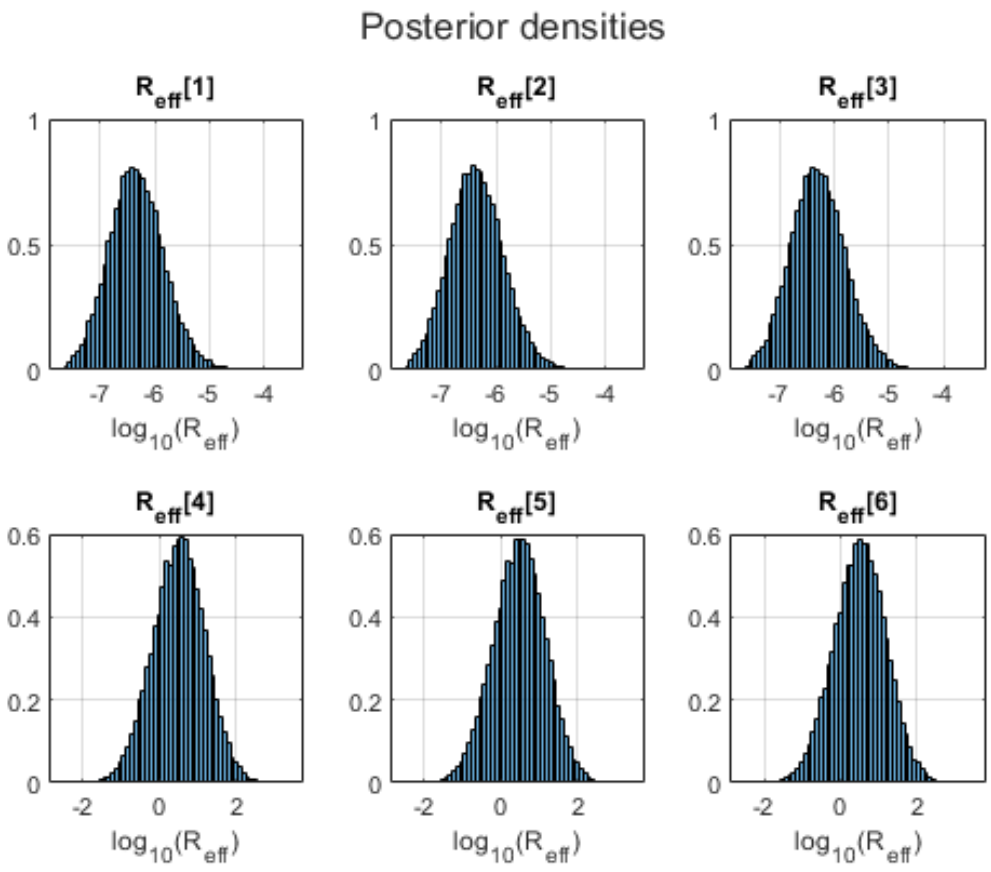


Figure 4: Posterior densities for the effective decoy-to-electronics energy ratio R_{eff} (three OpenBUGS chains pooled), plotted as $\log_{10}(R_{\text{eff}})$. Scenarios 1–3 concentrate near 10^{-7} (negligible decoy dominance), whereas scenarios 4–6 concentrate around $R_{\text{eff}} \sim 1$–10 with a long upper tail, indicating that decoy-dominated coupling is plausible but uncertain under current priors. Because the horizontal axis is $\log_{10}(R_{\text{eff}})$, values to the right of zero correspond to $R_{\text{eff}} > 1$, i.e. configurations in which more energy is dissipated in the decoy than in the protected enclosure. Values exceeding 1 on this scale ($R_{\text{eff}} > 10$) indicate an order-of- magnitude dominance of the decoy pathway, implying that the majority of coupled energy is intercepted externally rather than delivered to vulnerable electronics.

Posterior inference from the illustrative dataset reveals a regime transition governed primarily by coupling and resonance alignment. Baseline scenarios (1–3) concentrate at $\log_{10}(R_{\text{eff}}) \ll 0$, implying negligible decoy dominance. However, modest increases in coupling shift the posterior mass across $\log_{10}(R_{\text{eff}}) = 0$, with substantial probability extending beyond $\log_{10}(R_{\text{eff}}) = 1$, corresponding to order-of-magnitude decoy dominance. The electromagnetic analysis demonstrates that resonance placement is controlled principally by lateral cavity dimensions rather than depth for the dominant modes observed in the present simulations, providing a practical geometric tuning mechanism. Taken together,

these results indicate that the decoy-dominated regime is not a fragile artefact of priors but an engineerable operating point accessible through resonant design. The results remain illustrative and depend on the assumed calibration between simulated absorption features and effective coupling, motivating targeted experimental measurements to constrain the relative coupling coefficients and reduce posterior uncertainty.

The dissipation fractions η_e and η_{d0} remain weakly constrained and centred near their priors, indicating that the early-phase proxy data primarily inform relative coupling coefficients rather than intrinsic loss factors. The nonlinear enhancement parameter α is moderately identified, with posterior mass favouring positive values consistent with cumulative decoy dissipation.

The disruption threshold E_{th} remains broadly distributed, and the resulting failure probabilities cluster near 0.5. This behaviour is expected in early benchtop testing, where measurements constrain coupling pathways but do not directly observe electronic upset. Higher-energy or facility measurements would progressively constrain E_{th} and sharpen predicted disruption probabilities.

7.8. Linking Full-Wave RF Modelling to the Bayesian Coupling Model

The full-wave simulations presented in Figures 2–3 provide frequency-dependent absorption proxies defined as

$$A(f) = 1 - |S_{11}(f)|^2,$$

where $S_{11}(f)$ is the simulated reflection coefficient of the decoy structure. Although $A(f)$ does not directly represent energy delivered to protected electronics, it quantifies the strength of externally supported resonant coupling features. In the Bayesian framework introduced above this quantity informs the decoy coupling coefficient $\kappa_d(f)$ through a calibrated mapping

$$\kappa_d(f) = c_{\kappa,d}\, A(f),$$

where $c_{\kappa,d}$ is treated as an uncertain calibration constant with a prior distribution reflecting incomplete knowledge of real-platform coupling pathways. In this way, the full-wave model provides a physics-based prior shape for frequency- dependent decoy behaviour, while uncertainty in geometric placement, platform integration, and excitation conditions is explicitly propagated through the hierarchical model.

The emergence of multiple resolvable cavity resonances (e.g., the additional feature near 5.77 GHz in the 60 mm configuration) implies that $\kappa_d(f)$ may exhibit multi-modal structure rather than a single narrow peak. This modal richness expands the feasible design space and supports the hypothesis that engineered external resonances can create preferential coupling pathways without requiring access to high-power threat systems during early-stage evaluation.

As experimental measurements (e.g., VNA-derived scattering data or facility coupling tests) become available, the calibration constants and dissipation parameters can be updated within the same Bayesian framework, progressively tightening posterior uncertainty on the effective decoy-to-electronics energy ratio R_{eff}.

7.9. Scope and limitations

This minimal model intentionally compresses the EM interaction into a small number of identifiable parameters, enabling immediate Bayesian updating from early measurements. It does *not* attempt to resolve detailed cavity fields, platform aspect dependence, or explicit antenna coupling; those refinements can be introduced later by replacing the proxy mapping $\kappa \leftarrow A$ with a more explicit coupling sub-model. Nevertheless, the competing-pathway structure is faithful to the device concept and its intended mechanism: diversion of HPM energy into an external sacrificial resonance rather than into the protected enclosure.

8. Design Implications

The combined full-wave modelling and Bayesian inference results yield several concrete engineering implications for the practical development of sacrificial resonant decoys.

8.1. Electrical Size Determines Frequency Placement

Full-wave simulations of 50–60 mm cavity structures demonstrate that strong absorption-proxy peaks emerge when the lateral dimensions are of order half a wavelength in the frequency band of interest. For threat bands in the 4–6 GHz region, this corresponds to characteristic lateral dimensions of approximately 45–65 mm. The appearance of dominant resonances near these frequencies is consistent with classical rectangular cavity modal scaling.

This implies that decoys cannot be arbitrarily miniaturised: to present a strong engineered coupling feature, the lateral dimension must be electrically large enough (approximately $\lambda/2$) to support resonant modes in the threat band. Frequency placement is therefore primarily controlled by lateral geometry rather than thickness.

8.2. Depth Plays a Secondary Role

Parametric sweeps over cavity depth (40–60 mm) show only modest frequency shifts of the dominant resonance for a fixed lateral size. This indicates that resonance placement is governed principally by lateral dimensions and excitation geometry, with depth acting as a secondary tuning parameter.

From a design perspective, this is advantageous: the decoy can be substantially thinner than a full protective Faraday enclosure while still supporting strong resonant behaviour. Lightweight, externally mounted sacrificial structures are therefore feasible without requiring large enclosure volumes.

8.3. Modal Richness Enables Multi-Band Behaviour

For larger lateral dimensions (e.g., 60 mm), additional higher-order resonances appear within the 3–7 GHz band. The presence of multiple resolvable modes indicates that compact cavity geometries can exhibit modal richness rather than a single narrowband response.

This expands the design space: increasing lateral dimension increases modal density and permits the deliberate placement of multiple coupling features within the operational

band. Such behaviour may be exploited to realise multi-band or partially broadband decoy configurations.

8.4. Relative Coupling, Not Absolute Absorption, Governs Effectiveness

The Bayesian model highlights that device effectiveness depends on the relative coupling strength between the decoy and the vulnerable electronics pathway. The relevant metric is the effective energy ratio R_{eff}, which compares energy delivered to the decoy with that delivered to protected electronics.

Posterior inference shows that modest changes in assumed coupling coefficients can shift the system between two distinct regimes: (i) decoy-irrelevant ($R_{\text{eff}} \ll 1$), and (ii) decoy-dominant ($R_{\text{eff}} \gtrsim 1$). The practical design objective is therefore not maximal absorption in an absolute sense, but the creation of a resonant pathway that is electromagnetically more attractive than natural enclosure leakage modes.

8.5. Practical Design Targets

For illustrative threat bands in the few-GHz regime, the modelling suggests the following preliminary targets:

1. Lateral dimension: approximately 45–65 mm (order $\lambda/2$).
2. Thickness: can be reduced relative to lateral size, subject to structural and integration constraints.
3. Multi-mode behaviour: increasingly accessible for lateral dimensions near or above 60 mm.

These values are not definitive prescriptions but provide physically grounded starting points for prototype fabrication.

8.6. Measurement Priorities

Given the sensitivity of R_{eff} to relative coupling, our measurement work will prioritise:

1. Measurement of decoy scattering behaviour (S_{11} and derived $A(f)$) using benchtop VNA equipment.

2. Comparative coupling measurements to representative wiring or enclosure structures to constrain calibration constants in the Bayesian model.

Reducing uncertainty in these coupling parameters will directly tighten posterior predictions and determine whether a given geometric design operates in the decoy-dominant regime.

9. Conclusions

This work has examined the concept of an externally mounted sacrificial resonant decoy as a protective measure against high–power microwave (HPM) disruption of small uncrewed aerial vehicle (UAV) electronics. The central design objective is competitive resonant dominance: under either passive illumination or frequency–agile probing, incident HPM energy will preferentially couple into the most strongly dissipative resonant pathway available to the target. By introducing a deliberately engineered, lossy resonance external to the protected enclosure, the decoy seeks to bias that competition away from unintended enclosure leakage modes.

Full-wave electromagnetic modelling demonstrates that resonance placement is governed primarily by lateral cavity dimensions, providing a practical geometric tuning mechanism. Depth variations within the investigated range produce comparatively modest frequency shifts, indicating that compact externally mounted structures can be designed to align resonances within anticipated threat bands. These results establish that the decoy response is not incidental but controllable through geometry.

A compact probabilistic representation of the competing pathways was then developed, in which enclosure and decoy coupling coefficients and dissipation factors jointly determine the relative energy deposition. The resulting dominance ratio, R_{eff}, provides a physically interpretable criterion for effectiveness: values $R_{\text{eff}} > 1$ correspond to preferential decoy deposition, while $R_{\text{eff}} \gg 1$ indicate order-of-magnitude dominance of the sacrificial pathway.

An OpenBUGS implementation was presented to propagate uncertainty in RF coupling and threat parameters into predicted disruption outcomes while allowing incremental updating from benchtop measurements. Illustrative inference shows that modest changes in coupling assumptions can shift the system between enclosure-dominated and decoy-dominated regimes, confirming that the protective effect depends primarily on relative coupling strength rather than absolute absorption alone. In contrast, the absolute disruption threshold remains weakly identified in the early data regime, indicating that higher-energy or facility measurements are required to constrain functional upset probabilities.

Several broader conclusions follow. First, sacrificial resonant decoys represent a physically grounded complement to conventional shielding, particularly against those frequency-adaptive HPM systems that seek efficient coupling pathways. Second, resonance alignment and relative coupling strength—rather than broadband absorption—are the key design variables. Third, Bayesian inference provides a tractable and extensible framework for integrating heterogeneous RF measurements and progressively refining effectiveness estimates in a domain where full-scale adversarial testing is impractical.

Overall, the analysis indicates that externally mounted resonant decoys can, in principle, be engineered to dominate available coupling pathways and thereby reduce energy delivery to protected UAV electronics. The combined electromagnetic and probabilistic framework developed here supports rational geometry selection, targeted measurement, and structured uncertainty reduction for counter-HPM resilience.

References

[1] Jack Watling et al. *Stormbreak: Fighting Through Russian Defences in Ukraine's 2023 Offensive*. Royal United Services Institute (RUSI), 2023.

[2] International Institute for Strategic Studies. *The Military Balance 2024*. London: Routledge, 2024. DOI: 10.4324/9781003485834.

[3] Mark Hvizda et al. *Dispersed, Disguised, and Digitally Enabled: Observations from the Russia–Ukraine War*. RAND Corporation, 2025.

[4] Jakob Brenner and Matt Ferran. *L-297 Loitering Munitions*. NATO MSIAC, 2024.

[5] Kateryna Bondar. *Ukraine's Future Vision and Current Capabilities for Waging AI-Enabled Autonomous Warfare*. CSIS analysis. 2025.

[6] Graham V Weinberg. "Effect of dwell time, distance, and reflecting area on UAV swarm elimination in the presence of directed energy weapons". In: *The Journal of Defense Modeling and Simulation* 0.0 (0), p. 15485129251340992. DOI: 10.1177/15485129251340992. eprint: https://doi.org/10.1177/15485129251340992.

[7] Derek B Worth, Brian G Woolley, and Douglas D Hodson. "SwarmSim: a framework for modeling swarming unmanned aerial vehicles using Hardware-in-the-Loop". In: *The Journal of Defense Modeling and Simulation* 18.2 (2021), pp. 105–124. DOI: 10.1177/1548512917715154. eprint: https://doi.org/10.1177/1548512917715154.

[8] John Arquilla and David Ronfeldt. *Swarming and the Future of Conflict*. Tech. rep. MR-1100-OSD. RAND Corporation, 2000.

[9] Craig W. Reynolds. "Flocks, Herds, and Schools: A Distributed Behavioral Model". In: *Proceedings of SIGGRAPH '87*. ACM, 1987, pp. 25–34. DOI: 10.1145/37401.37406.

[10] Iain D. Couzin et al. "Collective Memory and Spatial Sorting in Animal Groups". In: *Journal of Theoretical Biology* 218.1 (2002), pp. 1–11. DOI: 10.1006/jtbi.2002.3065.

[11] Vitalii Savchenko et al. “Model of an alternative navigation system for high-precision weapons”. In: *The Journal of Defense Modeling and Simulation* 19.3 (2022), pp. 255–262. DOI: 10.1177/1548512920921955. eprint: https://doi.org/10.1177/1548512920921955.

[12] U.S. Government Accountability Office. *Directed Energy Weapons: DOD Should Focus on Transition Efforts to Ease Fielding Challenges*. Tech. rep. GAO-23-105868. 2023.

[13] Laura Middeldorp, Kerry Malone, and Wouter Noordkamp. “Quantifying the robustness of a Bayesian Belief Network in the context of Unmanned Aerial System threat prediction”. In: *The Journal of Defense Modeling and Simulation* 22.4 (2025), pp. 359–371. DOI: 10.1177/15485129231206825. eprint: https://doi.org/10.1177/15485129231206825.

[14] *High Power Microwave Threat to Infrastructure and Military Equipment*. NATO Science and Technology Organization (STO), 2010. DOI: 10.14339/RTO-TR-SCI-132.

[15] *Tactical Implications of High Power Microwaves*. NATO RTO, 2003. DOI: 10.14339/RTO-TM-028.

[16] Dale Reding et al. *Science & Technology Trends 2023–2043: Volume 2*. NATO, 2023.

[17] Zhao Zhang et al. “Strong Electromagnetic Interference and Protection in UAVs”. In: *Electronics* 13.2 (2024). ISSN: 2079-9292. DOI: 10.3390/electronics13020393.

[18] Air Force Research Laboratory. *AFRL conducts swarm technology demonstration (THOR)*. News release. 2023.

[19] Epirus Inc. *Leonidas High-Power Microwave Defeats 49-Drone Swarm (100% of Drones) at Live-Fire Demonstration*. Press release. 2025.

[20] William A. Radasky, Carl E. Baum, and Martin W. Wik. “Introduction to the Special Issue on Intentional Electromagnetic Interference (IEMI)”. In: *IEEE Transactions on Electromagnetic Compatibility* 46.3 (2004), pp. 314–321. DOI: 10.1109/TEMC.2004.833449.

[21] James Benford, John A. Swegle, and Edl Schamiloglu. *High Power Microwaves*. 3rd ed. Boca Raton, FL: CRC Press, 2016. ISBN: 978-1498722687.

[22] *IEC 61000-2-13: EMC – Part 2-13: Environment – High-power electromagnetic (HPEM) environments*. International Electrotechnical Commission, 2005.

[23] "High Power Microwave Weapon System". US20180058826A1. 2018.

[24] "Submicrosecond Long-Pulse, High-Efficiency Relativistic Cherenkov Oscillator". Patent CN106253031B. Chinese patent (granted). Also referenced as "submicrosecond-grade long pulse high-efficiency relativistic Cherenkov oscillator." Reviewed via project files. 2021.

[25] "High-Power Microwave Generating Device and Method". Patent CN111192804B. Chinese patent (granted). 2022.

[26] "High-Power W-Band Microwave Generation System". Utility Model CN216348091U. Chinese utility model. 2022.

[27] "High-Power Microwave Generating Device". Utility Model CN211654764U. Chinese utility model. 2020.

[28] "Device for damaging electronic equipment using unfocused high power millimeter wave beams". Patent US5777572A. Granted U.S. patent. Reviewed via project files. 1998.

[29] Guoqi Ni, Benqing Gao, and Junwei Lu. "Research on High Power Microwave Weapons". In: *Asia-Pacific Microwave Conference (APMC)*. 2005.

[30] Andy Lowery. *Microwave Weapons vs. Copper-Taped Drones*. YouTube video (Epirus / channel: The Merge). Quote at 03:08–03:56 of the video (transcript). "all of these drones require penetrations ... the holes actually are quite convenient for us because we can tune our back end...". 2024.

[31] David Lunn et al. "The BUGS project: Evolution, critique and future directions". In: *Statistics in Medicine* 28.25 (2009), pp. 3049–3067. DOI: 10.1002/sim.3680.

[32] Martyn Plummer. "JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling". In: *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*. Vienna, Austria, 2003.

[33] David J. Lunn et al. "WinBUGS – A Bayesian modelling framework: Concepts, structure, and extensibility". In: *Statistics and Computing* 10.4 (2000), pp. 325–337. DOI: 10.1023/A:1008929526011.

[34] Bob Carpenter et al. "Stan: A Probabilistic Programming Language". In: *Journal of Statistical Software* 76.1 (2017), pp. 1–32. DOI: 10.18637/jss.v076.i01.

[35] David J. Spiegelhalter et al. "Bayesian Analysis Using Gibbs Sampling: BUGS". In: *MRC Biostatistics Unit Technical Report* (1996).