

Holding Ground with Machines That Do Not Move

Why Endurance Is the Wrong Measure for Terrain Denial, and What a Thirty-Day Armed Emplacement Actually Costs in Watt-Hours

Christopher I. V. Farmer

Independent researcher, Geneva, Switzerland

6 August 2026

Suggested citation. C. I. V. Farmer, “Holding Ground with Machines That Do Not Move”, preprint, v1.0.0 (2026)..

Status. Preprint — not peer reviewed. Version v1.0.0.

Keywords. unmanned ground vehicles; terrain denial; emission control; unattended ground sensors; energy autonomy; closed systems; electronic warfare; military robotics

Licence. Creative Commons Attribution 4.0 International (CC BY 4.0).

Contents

Holding Ground with Machines That Do Not Move	2
Why Endurance Is the Wrong Measure for Terrain Denial, and What a Thirty-Day Armed Emplacement Actually Costs in Watt-Hours	2
Abstract	2
Note on sources and method	4
Part I—Introduction	5
1. The problem stated	5
2. Why the current recommendation is the wrong answer	5
3. What this paper does	6
4. What is new here, and what is not	6
Part II—Two Axes, Not One	9
5. The conflation	9
6. Axis one: closure, and the withdrawal test	10
7. Axis two: emission posture	11
Part III—The Dependency Audit	12
9. What a control stack actually consumes	12
10. What the audit does not establish	15
Part IV—The Emission Budget	16
11. What this Part can and cannot establish	16
12. The asymmetry that governs everything	16
13. Locating, which is a harder problem than detecting	18
14. The postures compared	18
15A. The reporting problem, and a rule for when to break silence	19
Part V—The Architecture Specified	19
16. A holding platform, described so that it can be argued with	19
Part VI—Energy	24
17. The discipline is measuring the wrong quantity	24
18. The architectural inversion	26
19. The interlock	30
20. Delivery rather than generation	32
21. Scavenging	33
21A. What energy buys: the discrimination exchange rate	34
22. What none of this solves	36
Part VII—What Closure Costs	37
23. The risks, stated at full strength	37
24. Which mitigations actually work	38
25. Receive-only, which is the most valuable design choice available	38
26. Expiry into inertness	40
26A. The adversary's counter-architecture	40
27. What remains unmitigated	41
27A. What this makes the system, doctrinally	42
Part VIII—The Boundary, and Conclusions	43
28. Where closure is the wrong answer	43
29. Conclusions on each element	44
30. What would change the analysis	46

31. Summative remarks	47
Appendix—What Would Have to Be Built	47
A.1 The experiment that matters most	48
A.2 The power measurements	48
A.3 The harvesting and cold tests	48
A.4 The scavenging feasibility study	49
A.4A The displacement-detection test	49
A.5 The dependency audit, repeated properly	49
A.6 What the paper got wrong in draft, recorded because it bears on the argument	50
A.7 Supplementary material	50

Holding Ground with Machines That Do Not Move

Why Endurance Is the Wrong Measure for Terrain Denial, and What a Thirty-Day Armed Emplacement Actually Costs in Watt-Hours

Christopher I. V. Farmer

Abstract

Unmanned ground vehicles have moved from experiment to mass employment in under three years. Ukraine logged nearly 24,500 ground-robot missions in the first quarter of 2026 and more than nine thousand in March alone, contracted 25,000 vehicles for the first half of the year, and states the objective that all frontline logistics be performed by robotic systems. In some brigades these platforms already perform 80 to 90 per cent of forward logistics. One of them, mounting a machine gun, held a fighting position for over a month.

That last case is the subject of this paper, and it does not fit the discipline that produced it. Fielded ground robots are rated at four to eight hours of endurance for small platforms and eight to sixteen for mid-sized logistics types, with a stated next-generation ambition of endurance measured in days. A month-long emplacement is two orders of magnitude beyond that figure. It happened anyway, and the reason it happened is that the figure measures the wrong thing.

This paper makes three arguments.

The first is definitional and it clears away a confusion that is doing real damage. “Offline” is used to mean two distinct things: that a system depends on no infrastructure a third party controls, and that a system does not emit. Only the first is an architectural property. The paper separates them into two axes and supplies a test for the first that any reader can apply to any system: *does the mission fail if a party other than the operator withdraws a service?* A system that passes is closed. Closure excludes control paths across commercial constellations, remote inference, vendor-delivered updates, licence and attestation servers, and telemetry to a manufacturer. It permits a private encrypted link to a named platform or group, pre-shared keys, a fibre tether, and inference resident on the platform. Closure and silence are then independent: a closed system may run at any emission posture from no link at all to a continuous one, and that choice is a mission setting rather than a commitment made at design time. The prevailing recommendation that ground robots be networked into wider strike architectures asks for openness on the first axis in order to obtain coordination, which is available on the second at no cost to closure.

The second is empirical and it is a dependency audit. For a representative fielded control stack the paper enumerates every party whose withdrawal breaks the mission—constellation operator, inference provider, firmware signer, licence server, map source, frequency coordinator—and counts them. The result is a measure of how far current practice sits from closure, and it is further than the discussion assumes.

The third concerns energy, and the paper’s principal contribution because energy is the binding constraint on holding ground. Endurance is quoted in drive hours. Locomotion dominates the budget of a moving platform and is nearly absent from that of a static one, so the published figures say almost nothing about a mission whose defining feature is that the platform does not move. The paper proposes the quantity the mission actually needs—**watt-hours per emplaced-day at a stated readiness level**—and above it a formation-level measure, **energy per square kilometre per day of denial**, which permits robotic and infantry terrain denial to be compared on a single axis. Holding ground has always been an energy problem; robotic holding is the first form of it that is directly measurable.

Four routes to a sustainable budget follow. The largest is an architectural inversion borrowed from an adjacent field: unattended ground sensors achieve multi-month emplacements because they are event-driven, idling at microwatt levels between scan events and waking through tiers as a candidate develops. Ground robots are not built this way, because they are designed as vehicles that happen to stop rather than as sensors that happen to have effectors. The paper argues the holding platform should be architected on the sensor model, and that this changes the required budget by orders of magnitude rather than by a factor.

A second result follows from the first and is the paper’s most consequential claim. Ambient harvesting is normally dismissed for this application at high latitude in winter, on yield and on signature. That dismissal assumes a vehicle-class load. Against a microwatt floor the arithmetic changes, and a small terrain-matched harvesting surface becomes sufficient rather than marginal. Closure reduces emission; reduced emission collapses the duty cycle; the collapse makes harvesting viable; viability removes the resupply dependency, which is itself closure. The two axes and the energy budget are not three problems but one.

The remaining two routes are delivery and scavenging. For loads harvesting cannot meet, the answer is battery **exchange** rather than recharge, performed by the autonomous logistics fleet already conducting most forward resupply in some formations: it takes minutes rather than hours, requires the holding unit neither to move nor to emit, and separates generation, which can occur anywhere, from delivery, which is contested and already robotic. Scavenging is the speculative route and is labelled as such throughout: contested ground is dense with stored energy, and expended first-person-view drones—produced at a combined rate exceeding ten million units a year and each carrying a lithium-polymer pack—constitute a supply distributed across precisely the terrain a holding system occupies, and one that grows as fighting intensifies.

Read as a price list rather than as a constraint, that budget yields the paper’s most portable result. Identification and release judgement account for 91 per cent of it; recallability for 1.3 per cent; the weapon itself for 0.5 per cent. A system that denies ground without discriminating costs 81 mWh per emplaced-day and one that looks first costs eleven times that—thirteen on the paper’s own improved receiver, so the figure quoted is the conservative one—so **what the energy purchases is discrimination**. A mine is the zero-energy limit of the same design: it spends nothing and exercises no judgement. The two are the same class of object separated by a measurable expenditure, which makes the question of how much discrimination should be required in such systems answerable, in this class at least, as a quantity of watt-hours. That question is live rather than academic. Amended Protocol II already imposes self-deactivation, detectability and a thirty-day horizon on the zero-energy case; on the figures here, the first two cost approximately nothing to meet in the energy-spending case and discrimination—the obligation the law does not impose on mines, because no device without sensing could meet it—costs 830 mWh per emplaced-day. The reason mines are regulated by marking, mapping and self-destruction rather than by a duty to discriminate is that discrimination was unavailable at any price. It is now available, and this paper prices it.

Two further results follow from the same analysis. A holding system’s intelligence product competes with its survivability for one account, and the paper states the rule that decides between them and the design consequences of setting reporting thresholds at emplacement. And against a competent adversary the architecture resists counter-measure structurally rather than technically—jamming a receive-only system denies recall without preventing function, so it makes the system more dangerous rather than less—with the consequence that its real vulnerability lies not in the emplaced platform but in the emplacement, exchange and recovery cycle, and that the value of the energy work is measured in how rarely the system must be visited rather than in how cheaply it runs.

The escalation rate that governs the whole budget is not left as an assumption. It is the output of a detection

cascade whose terms are separately measurable, and writing it out yields a result neither parent literature contains: **the energy cost of holding ground is set by the false-alarm rate rather than by enemy activity**, because environmental false trips greatly outnumber genuine events in any unattended sensor. A sector crossed daily costs almost what an empty one costs; a sector in high wind or beside a road costs several times more. Two design consequences follow—that false-alarm rejection in the cheap tiers is the highest-leverage parameter available, and that detection range and energy are coupled, so the optimum sensitivity sits below what the hardware permits.

Three results follow from the platform being static, which the field has not examined because it designs for motion. A static sensor subtracts its background from a single survey and so needs no continuous ranging, removing lidar from both the emission and the power budget. A static platform’s cost is what it pays to do nothing, which inverts the ordering of morphologies and makes the legged designs now attracting investment the worst choice for this mission. And the static frame conceals a defect that undermines the paper’s own safety argument: a platform displaced by blast, subsidence or a vehicle strike continues to enforce its surveyed sector over ground nobody surveyed. The answer is self-localisation by inter-feature geometry—pairwise ranges and angles among fixed features are invariant under displacement of the observer where absolute bearings are not—which distinguishes a moved platform from a changed world using no external reference at all, and therefore preserves closure where positioning or inertial aids would not.

One further result is doctrinal rather than technical. Countermobility doctrine plans emplaced effects by intent, and every obstacle in that taxonomy has its intent fixed at the moment it is laid. A closed system at posture 1 is the first that can be re-tasked afterwards—narrowed, relaxed, made inert, revived—by a command costing about one part in a thousand of its energy budget and radiated from the rear rather than by the obstacle itself. Gaps and lanes become parameter changes rather than engineer tasks conducted under fire. That is the strongest argument for building the recall channel, and it is not a safety argument: without the channel the object is a mine with better sensors.

Two fielded systems bound the argument at either end. Anduril’s autonomous surveillance towers, deployed past three hundred units and classifying at thousands of metres, show what range buys and what it costs: height, continuous power, active sensing, and a design whose stated value is the network it feeds—dependent by design, and correctly so for the customer it was built for. Ukraine’s Sky Sentinel, a combat-tested static armed turret that detects, tracks and fires without an operator, shows where this paper’s architecture stops: it is radar-cued because its targets close at hundreds of miles an hour. **The speed of the target sets the sensing modality, and the modality sets the emission and the power budget.** Everything here is available for ground denial because ground targets are slow.

The paper states what each proposal would require to be validated, and marks every figure that is derived rather than measured. Its central engineering claim has not been tested on hardware. What it offers is a specification, an arithmetic, and the experiment that would decide it.

Keywords: unmanned ground vehicles; terrain denial; emission control; unattended ground sensors; energy autonomy; closed systems; electronic warfare; military robotics

Note on sources and method

This paper draws on three classes of source and treats them differently, because they do not carry the same weight and presenting them as though they did would be the principal way to go wrong.

Measured and published engineering values—power figures, energy densities, photovoltaic yields, propagation and detection relationships—are taken from the technical literature and from manufacturers’ published data, and are cited to the record. Where a value is reported as a range the range is given rather than its midpoint.

Field employment data is taken from operational reporting, from statements by the ministries and general staffs concerned, and from defence-technical journalism. It is the least controlled class of evidence here.

Figures issued by a party to a conflict about its own performance are attributed to that party in the text rather than stated as fact, and where two independent reports differ the difference is shown.

Derived quantities are the author’s own arithmetic from the two classes above. Every one is marked, its inputs are given, and the derivation is reproducible from the stated figures. Where a derivation rests on an assumption that cannot be checked, the assumption is stated at the point of use and its effect on the result is shown.

Every authority has been resolved against a primary or authoritative record before citation. Anything that could not be verified was removed rather than approximated.

The distinction that matters most in this paper is between what closes arithmetically and what has been demonstrated. The architecture proposed in Part VI has not been built. Its power budget closes on published component figures; that is a different and much weaker statement than a bench result, and it is not offered as one. The Appendix sets out what would have to be built and measured to move each claim from the first category to the second.

Part I—Introduction

1. The problem stated

A machine gun mounted on a tracked chassis held a fighting position on the Ukrainian front for more than a month.¹ It was operated by the 3rd Assault Brigade, it shielded infantry from repeated attacks, and it did the thing infantry are held in place to do: it made ground expensive to cross.

Nothing in the published engineering literature explains how that was possible. Fielded unmanned ground vehicles are rated at four to eight hours of endurance for small platforms and eight to sixteen hours for mid-sized logistics types.² The stated ambition for the next generation is endurance measured in days rather than hours.³ A month is between forty and two hundred times the quoted figure depending on the platform compared: 720 hours against a sixteen-hour upper bound is forty-five, and against a four-hour lower bound is a hundred and eighty. The discrepancy is not a measurement error and it is not a claim that should be doubted. It is a sign that the quoted figure measures something other than the mission.

The figure measures **drive hours**. Locomotion dominates the energy budget of a moving ground platform and is nearly absent from that of a static one. A system emplaced to deny ground spends its energy on sensing, on computation, on standby, and briefly on effect. It spends almost none on moving, because not moving is the mission. The discipline has no metric for that case, which is why it has no design method for it either.

This paper is about the case. It asks what a system built specifically to hold ground should look like, what it should depend on, what it should emit, and how it should be powered for a month. The three questions turn out to be one question.

2. Why the current recommendation is the wrong answer

The prevailing view is explicit and recent. Assessing Ukraine’s ground robots in March 2026, the Modern War Institute concluded that the lesson for Western forces is to build for networking: ground robots as nodes in a robotic strike group, drawing navigation and targeting data from overhead quadcopters, tied into existing drone, electronic-warfare and logistics systems, and that platforms which do not plug into those networks

¹*Militarnyi*, ‘UGV with a machine gun replaced infantry and held a position for a month and a half’ (*Militarnyi*, 2025) <https://militarnyi.com/en/news/ugv-with-a-machine-gun-replaced-infantry-and-held-a-position-for-a-month-and-half/> accessed 5 August 2026. The platform was operated by the 3rd Assault Brigade. The report is Ukrainian and concerns Ukrainian performance; it is treated here as a claim by a party to the conflict, and the argument of this paper does not depend on the precise duration.

²Unmanned Systems UK, ‘UGV Power Systems and Energy Management’ (15 June 2025) <https://unmanned-systems.co.uk/blog/2025-06-15-ugv-power-systems-energy-management/> accessed 5 August 2026, reporting four to eight hours for small inspection platforms and eight to sixteen for mid-sized logistics types.

³*ibid*.

“will sit in storage yards.”⁴ The recommendation is careful, it is drawn from operational experience, and its premise is correct: a platform operating alone in this environment is at a severe disadvantage.

The conclusion does not follow from the premise, and the reason is a conflation this paper’s Part II exists to remove. The argument for networking is an argument for **coordination**—shared picture, cued detection, mutual support, handoff between platforms. It is not an argument for **dependency**. Those are separable, and the current generation of systems has bundled them because the cheapest way to obtain coordination happens to be to buy it from a third party.

The cost of that bundle is not hypothetical. Ukrainian ground robots have been controlled over a commercial satellite constellation whose availability has been reported as a matter of private decision, and whose measured throughput to these platforms has been reported as low as 10 Mbit/s with video quality poor enough to impair use.⁵ That is a control path with a proprietor. Alongside it sit vendor-delivered software updates, remote licence checks, and map and terrain data drawn from external services. Each is a party who can withdraw.

A private encrypted link across a controlled group of platforms delivers the coordination the networking argument wants and introduces no such party. The question is therefore not whether to communicate. It is what to communicate with, and whose permission the mission requires.

3. What this paper does

Part II separates the two axes and states the test for the first. It defines closure by whether the mission survives the withdrawal of a service by a party other than the operator, sets out the emission ladder as a separate and orthogonal variable, and shows that treating them as one variable is what produces the false choice between capability and survivability.

Part III is a dependency audit of a representative fielded control stack, enumerating and counting the parties whose withdrawal breaks the mission. It is the paper’s first empirical contribution and it establishes how far current practice sits from closure.

Part IV quantifies the emission budget. For each posture on the ladder it states what is radiated, in what band, at what duty cycle, and what that implies for time-to-detection and time-to-geolocation against a competent effort. This is the part most exposed to a knowledgeable reader and it is explicit about which relationships are measured and which are modelled.

Part V specifies the architecture: sensing, on-platform inference, the authority envelope, key handling with no network at any point in the life cycle, release interlocks, and the sustainment concept.

Part VI is the paper’s principal contribution and concerns energy. It proposes the metric the mission requires, argues for an architectural inversion drawn from unattended ground sensors, shows the interlock between emission posture and harvesting viability, sets out battery exchange as the delivery model, and examines battlefield scavenging as the speculative route.

Part VII takes the risks at full strength: what is lost when a system cannot be reached, what a private link is worth when the platform can be captured, and what closure removes that supervision would otherwise have supplied. Mitigations are assessed for what they actually close.

Part VIII states the boundary—where closure is the wrong answer—and concludes.

4. What is new here, and what is not

A paper of this kind should say plainly which of its claims are original and which are restatements. The distinction is easy to blur and expensive to get wrong.

⁴Jorge Rivero, ‘Networked for War: Lessons from Ukraine’s Ground Robots’ (*Modern War Institute*, 9 March 2026) <https://mwi.westpoint.edu/networked-for-war-lessons-from-ukraines-ground-robots/> accessed 5 August 2026.

⁵ibid, reporting Starlink data rates to ground robots as low as 10 Mbit/s with consequent degradation of video quality.

4.1 What is not new

That electromagnetic emission is the principal route to detection, and that controlling it is a primary tactical concern, is settled doctrine. Emission control is a mature discipline with a substantial literature, and its central proposition—that the most effective means of evading signals intelligence is not to emit during periods of acute threat—is stated in the doctrinal sources and is not this paper’s discovery.⁶ The recent work on emissions control standard operating procedures for ground formations is likewise established practice being written down, not new theory.⁷

That event-driven sensor architectures achieve very long unattended deployments is the founding result of the unattended ground sensor field and dates to its earliest systems.⁸ Seismic and acoustic front ends idling at microwatt levels between scan events are ordinary engineering in that discipline. Nothing about the technique is claimed here.

That dependency on commercial space infrastructure is a strategic vulnerability has been argued at length since well before the present conflict, and the specific case has been examined by others.⁹

That battery exchange is faster than recharge, and that robotic logistics can carry energy as readily as it carries ammunition, are obvious once stated and have been implemented in commercial robotics.¹⁰

Priority is not claimed for any of these.

4.2 What is offered as new

First, the two-axis separation and the withdrawal test (Part II). The discussion of autonomy in ground systems runs on a single axis from “tethered to a human” to “independent”, and it treats connectivity as the same variable as dependency. Separating them, and supplying a test for dependency that can be applied to a specification rather than argued about, is offered as new. Its value is that it dissolves the choice the field believes it faces: coordination and closure are not in tension, and the systems that appear to demonstrate a tension are systems that bought coordination from a proprietor.

Second, the dependency audit as a method and as a result (Part III). Enumerating the parties whose withdrawal breaks a mission is not a difficult idea, but it does not appear to have been done for this class of system, and the count is larger than the discussion assumes.

Third, the emplaced-day metric and the formation-level denial quantity (Part VI §17). Endurance in drive hours is the wrong measure for a static mission, and no published measure exists for the right one. Watt-hours per emplaced-day at a stated readiness level is proposed as the platform quantity, and energy per square kilometre per day of denial as the formation quantity. The second permits robotic and human terrain denial to be compared on one axis, which has not been possible.

Fourth, and this is the central engineering claim, the architectural inversion (Part VI §18). Ground robots are designed as vehicles. Unattended ground sensors are designed as event-driven sensors. For the holding mission the second is the correct architecture and the first is not, and the two literatures do not cite one another. Building an armed static platform on the sensor power model is the proposal, and the claim is that it changes the achievable emplacement by orders of magnitude rather than by a factor.

Fifth, the interlock (Part VI §19). The same design decision that reduces detectability makes ambient harvesting viable, because harvesting is dismissed against a load that the decision removes. If this holds it is the most useful result in the paper, and §19 states the arithmetic and the assumption it rests on.

Sixth, energy as the price of discrimination (Part VI §21A). The budget re-read as a price list rather than as a constraint, with a mine as the zero-energy limit of the same design and the obligations of Amended Protocol II priced against it. This is the paper’s most portable result and it is not an engineering

⁶See Part IV §12 below for the doctrinal sources and their limits.

⁷US Army, ‘Adapting to Multi-Domain Battlefield: Developing Emissions Control SOP’ https://www.army.mil/article/284546/adapting_to_multi_domain_battlefield_developing_emissions_control_sop accessed 5 August 2026.

⁸See Part VI §18 for the unattended ground sensor literature.

⁹See Part III §9.

¹⁰See Part VI §20.

result: it makes the question of how much discrimination a system of this class should exercise answerable as a quantity of watt-hours, which is a question the legal literature currently argues without a unit.

Seventh, the reporting rule (Part IV §15A). That a holding system’s intelligence product and its survivability draw on one account is obvious once stated and is nowhere stated, because the literature treats links as control channels rather than as reporting channels. The rule for deciding between them, and the consequence that reporting thresholds belong in the emplacement act alongside the authority envelope, are offered as new.

Eighth, the design consequences of stasis, which the field has not worked out because it designs for motion (Part V §16.2A and Part VI §17.1A). Three follow and none appears in either parent literature. A static sensor subtracts its background from a survey performed once, so it needs no simultaneous localisation and mapping, no ego-motion estimation and no continuous active ranging—which removes lidar from both the emission budget and the power budget and is part of why the arithmetic at §18.2 closes at all. The relevant question for a static platform is not what it costs to move but what it costs to do nothing, and on that measure the ordering of morphologies inverts: a tracked platform at rest draws nothing from its drivetrain while a legged one must lie down, latch mechanically, or pay to stand. And the static frame conceals a defect: a displaced platform still enforces its surveyed sector mechanically over ground nobody surveyed. The answer offered is self-localisation by inter-feature geometry, which is invariant under displacement of the observer, requires no external reference and therefore preserves closure—and which must gate weapon release rather than silently re-deriving a sector.

Ninth, the detection cascade as the source of the energy budget (Part VI §18.3). That the escalation rate governing the whole budget is an output of a measurable cascade rather than an assumption, that false-alarm leakage rather than enemy activity dominates it, and that detection range and energy cost are therefore coupled with an optimum sensitivity below the hardware maximum.

Tenth, the mutable obstacle (Part VII §27A). Countermobility doctrine plans emplaced effects by intent—disrupt, turn, fix, block—and every obstacle in that taxonomy has its intent fixed at emplacement, so universally that doctrine does not state it. A closed holding system at posture 1 is the first obstacle whose intent can be changed after emplacement without anyone going forward. This is offered as the paper’s doctrinal contribution, and it converts the recall channel from a safety feature into a capability a commander wants.

Eleventh, the counter-architecture analysis (Part VII §26A). That jamming a receive-only holding system makes it more dangerous rather than less, and that the architecture’s real vulnerability is the emplacement and exchange cycle rather than the emplaced platform, invert the usual optimisation and are the practical consequence most likely to be acted on.

Twelfth, the scavenging proposal (Part VI §21). Offered as speculative and labelled so throughout. What is new is the identification of expended first-person-view airframes as a dense, renewing, geographically co-located energy source, together with a statement of the four engineering problems that stand between the observation and a usable system.

4.3 What would falsify each claim

The **two-axis separation** fails if a coordination function of genuine operational value can be shown to be unobtainable without a third-party service. The paper’s position is that none has been identified; a counter-example would defeat it.

The **inversion** fails on measurement. If a duty-cycled sensing architecture adequate for reliable detection and classification at the required ranges cannot be operated within the power envelope §18 states, the claim is wrong. That is a bench question and §18 specifies the test.

The **interlock** fails if the harvesting surface required to meet the reduced load is large enough to reintroduce the signature problem it was supposed to avoid, or if yield at the stated latitude and season falls below the stated floor. Both are measurable and §19 gives the thresholds.

The **scavenging proposal** fails if the recoverable fraction of expended airframes is materially below the figure §21 assumes, or if state-of-health assessment on unknown damaged cells cannot be performed within the platform’s own energy and mass budget.

The **energy-as-price-of-discrimination reframing** fails if 830 mWh per emplaced-day does not in fact buy identification good enough to count as discrimination. Pricing an obligation is only useful if the thing priced is the thing required, and a cheap identification that is unreliable is not a cheap obligation but an unmet one. That is a performance question about the imaging tier and it is measurable.

The **reporting rule** fails if the expected value of an observation cannot be estimated before emplacement with enough confidence to set a threshold. The rule would then be correct and unusable, which for a design rule is the same as being wrong.

The **detection cascade** fails if F and q cannot be separated in practice, or—more seriously—if tier-two rejection cannot be improved without a corresponding loss of true detections. The argument assumes the two are separable, which is to say that the tiers’ receiver operating characteristics are good enough that q can be lowered while P is held. A poor characteristic would couple them and the highest-leverage parameter would cease to be a lever.

The **mutable obstacle** fails on command delivery. If a receive-only channel cannot be relied on under jamming often enough to plan around, the obstacle is mutable in principle and fixed in practice, and the doctrinal claim reduces to a statement about ideal conditions.

The **counter-architecture analysis** fails if a cheap wide-area method for locating concealed static emplacements exists or emerges—persistent overhead change detection, multispectral or thermal survey at scale, or anything else that makes finding them routine rather than laborious. The whole argument that clearance is expensive rests on detection being the hard part, and that is a contingent fact about current sensing rather than a permanent one.

The **displacement-detection method** fails on ground without static features to survey, where there is no pattern to hold invariant, and it fails more seriously if inter-feature geometry cannot be recovered reliably enough from the tier-three imagery already being captured—in which case detecting displacement would need a sensor and a power budget of its own and would stop being nearly free.

The **morphology preference** fails if a legged platform can latch its posture at negligible power. The claim would then reduce to a cost-of-transport point, which is already known and is not this paper’s.

The **emplaced-day metric** cannot be falsified, because it is a definition rather than a claim. It can be shown to be useless, which would follow if designers found it did not discriminate between architectures. §17 states what it should discriminate between.

Part II—Two Axes, Not One

5. The conflation

This section aims to separate two properties which the discussion of unmanned ground systems treats as one, and to supply a test for the first of them which can be applied to a specification rather than argued about. The separation is the precondition for everything that follows, because the design choices in Part V and the energy argument in Part VI are only available once it is made.

The word carrying the confusion is “offline”. It is used to mean that a system depends on no infrastructure outside the operator’s control, and it is used to mean that a system does not transmit. Those are different properties, they vary independently, and treating them as a single variable produces two errors in opposite directions.

The first error is to conclude that a system which must communicate cannot be independent. That is false, and every special reconnaissance patrol and every submarine on station is a standing demonstration of its falsity. A patrol carries self-contained equipment, operates on pre-briefed intent, holds its own cryptographic material, and communicates on a private link in windows of its own choosing. A ballistic missile submarine

sustains months of dwell at near-zero emission with complete operational capability and no dependency on anyone else’s network.¹¹ Neither is disconnected. Both are independent.

The second error is to conclude that a system which is independent must be silent. That is equally false and the error this paper’s own framing risked. Independence is a statement about whose infrastructure the mission requires. Silence is a statement about what the system radiates in a given period. A closed system can be operated at any point on the emission ladder, and moving along that ladder is a tactical decision taken during the mission rather than an architectural commitment taken at design.

6. Axis one: closure, and the withdrawal test

6.1 The test

A system is **closed** if, and only if, the following holds.

No party other than the operator can cause the mission to fail by withdrawing a service.

The test is applied to a specification, not to a running system, and it is applied by enumeration: list every service the mission consumes, identify who supplies each, and ask what happens to the mission if that supplier stops. A single service supplied by a party outside the operator’s control, whose loss defeats the mission, is sufficient to fail the test.

Three features of the formulation are deliberate.

It says **withdrawing**, not attacking. Denial by an adversary is a different problem with a different treatment, and conflating the two obscures the point. The concern here is the party who supplies the service in the ordinary course and who may cease to do so—for commercial reasons, under legal compulsion, under political direction, on a change of ownership, or because a contract lapsed. Those causes are not adversarial and no amount of hardening addresses them.

It says **the mission**, not the platform. A service whose loss degrades performance without defeating the mission does not fail the test. This matters, because it permits a system to use an external service opportunistically while remaining closed, provided the mission is specified to complete without it. Opportunistic use is not dependency. Designed reliance is.

It says **the operator**, not the state. A national force operating a platform under a licence from a foreign vendor is not the supplier of its own firmware, and the test is not satisfied by the vendor being an ally. Alliance is a relationship between governments and it does not follow the corporate ownership of a supplier through a change of control.

6.2 What the test excludes

Applying it produces a list which is longer than the usual discussion of “connectivity” suggests, because the usual discussion attends to the control link and stops.

Control paths across commercially operated constellations. The supplier is a company. Its terms, its availability and its willingness are its own.

Inference performed off the platform. A classification or targeting decision computed in a data centre is a service supplied by whoever runs the data centre, over a link supplied by whoever runs the link.

Vendor-delivered software updates. A platform which requires them to remain current is dependent on the vendor continuing to deliver them, and on the delivery channel remaining open.

Licence, entitlement and attestation checks. A platform which will not operate, or which degrades, when it cannot reach a server to confirm that it is authorised, has handed a switch to the holder of the server. This is the most severe case in the list because it converts a commercial mechanism into an operational one.

¹¹The submarine case is the closest existing analogue to a month-scale emplacement and is used throughout this paper as a reference architecture rather than as an authority. See US Naval Institute, ‘Mission Command Means Emissions Control’ (*Proceedings*, May 2021) vol 147/5/1,419 <https://www.usni.org/magazines/proceedings/2021/may/mission-command-means-emissions-control> accessed 5 August 2026.

Externally supplied map, terrain and positioning data, where the mission requires it to be current. Cached data held on the platform is not a dependency; a live feed is.

Frequency allocation and deconfliction supplied by another party, where operation without it is prohibited or unworkable.

Telemetry obligations. A platform which reports its status, position or usage to a manufacturer is disclosing to a party who did not need to be told, and the disclosure is itself a channel.

6.3 What the test permits

The list of what closure allows is what makes the concept useful rather than ascetic.

A private encrypted link between the operator and a named platform. The operator owns both ends, holds the keys, and chooses when it is used. No third party supplies anything.

A private mesh across a controlled group of platforms. This is the important one, because it is where the coordination the networking argument wants actually lives. Platforms sharing detections, cueing one another, handing off tracks and maintaining a common local picture do all of that within a perimeter the operator controls. There is no proprietor.

A fibre-optic tether. Physically laid, unjammable, and emitting nothing. It has its own vulnerabilities, treated in Part IV, but dependency is not among them.

Pre-shared cryptographic material, loaded before deployment by a ceremony the operator conducts.

Inference resident on the platform. Computation performed locally consumes energy, which Part VI addresses, and it consumes nobody's permission.

Cached external data, including maps, terrain models and signature libraries, loaded before deployment and used thereafter without reference to a source.

Opportunistic use of any external service, provided the specification does not require it.

6.4 Why this is the right line

The objection to the test is that it is arbitrary: every system depends on something, and a supplier who provides ammunition can withdraw as readily as one who provides bandwidth.

The distinction is real and it is temporal. A logistics dependency is discharged **before** the mission and can be verified before commitment: a commander who has the ammunition has it, and the supplier's later change of mind does not reach into an engagement already under way. A service dependency is discharged **during** the mission and cannot be verified in advance, because the thing being relied on has not happened yet. The relevant property is not whether a party can affect you but whether they can affect you *after you are committed and while you cannot substitute*. That is what the withdrawal test isolates, and it is why cached data passes and a live feed does not.

7. Axis two: emission posture

7.1 The ladder

Emission posture is a separate variable and it runs as a spectrum. Five postures are useful for design and they are stated here in the order of increasing radiated energy.

Posture 0—no link. The platform neither transmits nor receives. It executes on intent loaded before emplacement. Detectability from the electromagnetic spectrum is zero; recall is zero.

Posture 1—receive only. The platform listens and never transmits. It can be commanded, updated within limits, and told to stop. Because it does not radiate, it cannot be located by direction-finding of its own emissions. This posture is treated at length in Part VII because it recovers most of what posture 0 gives up at very little cost, and the field has not made enough of it.

Posture 2—duty-cycled two-way. The platform transmits in scheduled or event-triggered windows and is otherwise silent. Detectability is a function of window length, frequency and waveform rather than of the mere existence of a link.

Posture 3—continuous two-way. A persistent radio link. Full supervision, full interactivity, and a continuous emission to be found.

Posture 4—tethered. A physical fibre link. Continuous interactivity with no radiated emission, at the cost of a physical line to the platform’s position and a limit on range and mobility.

7.2 Why the ladder is orthogonal to closure

Each posture can be implemented in a closed or an open form, and the pairing is free. Posture 3 over a private link with pre-shared keys is closed; posture 3 over a commercial constellation is not. Posture 1 receiving vendor-signed updates from a public channel is open despite being nearly silent; posture 1 receiving commands on an operator-held key is closed. The postures describe how much a system radiates. Closure describes whose permission it needs. Neither implies the other.

The practical consequence is the finding this Part exists to produce: **the choice presented to the field as a choice between capability and survivability is not one.** The capability being defended is coordination, and coordination is available at posture 2 or 3 on a private mesh, which is closed. The survivability being defended is low detectability, and that is a matter of posture rather than of architecture. A commander may move down the ladder as threat rises and back up as it falls, on a platform whose closure never changes. Building the choice into the architecture, as the current generation does by acquiring coordination from a proprietor, forfeits that flexibility permanently in exchange for a capability that was obtainable without it.

The practical payoff is stated once here and not repeated. Separating the axes makes the dependency audit at Part III possible, because that audit is the withdrawal test applied by enumeration; it makes emission a design variable rather than a consequence, so that asking how much a system may emit no longer determines what it depends on; and it makes the energy argument at Part VI available, because the largest term in a holding platform’s budget is set by posture and duty cycle rather than by capability.

Part III—The Dependency Audit

9. What a control stack actually consumes

This section applies the withdrawal test of Part II §6 by enumeration to a representative fielded control stack, and reports the count. Our purpose is to establish how far current practice sits from closure, because the discussion of connectivity in this field attends almost exclusively to the control link and treats everything else as infrastructure that is simply there.

The method is stated first so that a reader can repeat it against a different system. For each service the mission consumes: name the service, name the party who supplies it, state whether the mission fails without it, and state whether the loss can be substituted **during** the mission. A service which fails the third question and cannot be substituted is a dependency in the sense of §6.1. A service consumed opportunistically, whose loss degrades but does not defeat, is not.

9.1 The stack

The system audited is a mid-sized armed or logistics ground platform of the type in general service in 2025 and 2026, controlled from a position several kilometres to the rear. It is a composite rather than a specific product, assembled from the published description of the control methods in use: satellite links, line-of-sight radio, fibre-optic cable and aerial relay nodes, selected by range, terrain and electronic warfare conditions.¹² Nothing turns on the composite being any particular vendor’s platform, and the audit is more useful for not being.

¹²Rivero (n 4), describing control by satellite link, line-of-sight radio, fibre-optic cable and aerial relay according to range, terrain and electronic warfare conditions.

Eleven classes of service are consumed. They are taken in order of severity, which is not the order in which they are usually discussed.

One—the satellite control path. Where the link runs across a commercially operated constellation, the supplier is a company, and the case is not hypothetical in any respect. Service was activated in Ukraine in February 2022 by the operator’s decision.¹³ The operator has used geofencing to make the service available in some territories and not others, so that connection has failed where coverage was withheld.¹⁴ A request to extend coverage over occupied Crimea during an operation in September 2022 was declined by the company’s principal.¹⁵ And the operator has specifically limited use of the service for military purposes **including drone control**, maintaining that the network was to remain civilian, before establishing a separate product line for government customers.¹⁶ The academic treatment of the episode is now substantial and reaches in general the conclusion this audit reaches in a particular case: private ownership of infrastructure on which a belligerent depends transfers a measure of decision to the owner.¹⁷

That is the withdrawal test failing in the strongest available form. The service was not withheld by an adversary. It was withheld by the supplier, on the supplier’s own reasoning, and in one instance the restricted use case was the exact one this paper concerns.

Two—licence, entitlement and attestation. A platform which will not operate, or which operates in a reduced mode, when it cannot reach a server to confirm that it is authorised has given a switch to the holder of that server. This is the most severe class in the audit because the mechanism is commercial in origin and operational in effect, and because it is often invisible in a specification that describes only the control link. It is also the class on which the least is published, which is itself a finding: a purchaser cannot audit what a vendor does not disclose.

Three—firmware signing and delivered updates. Where a platform requires vendor-signed images to remain current, and particularly where it requires them to counter an evolving electronic warfare threat, the vendor is supplying a service throughout the mission’s life. Ukrainian practice is instructive on the alternative: brigade workshops of ten to twelve engineers change frequencies and software configurations in response to new threats in real time, without routing through a distant depot or manufacturer.¹⁸ That is a closure decision taken for operational reasons under pressure, and the correct one.

Four—off-platform inference. A classification or targeting decision computed elsewhere consumes both the compute service and the link that reaches it. The dependency is doubled and the failure modes compound.

Five—global navigation satellite positioning. The supplier is a state. Withdrawal in the sense used here—a policy decision to degrade or deny regionally—is distinct from adversarial jamming and spoofing, which are treated in Part IV, and both are live. The mitigation is well understood and is itself a closure measure: inertial navigation with terrain-referenced or visual correction, held on the platform.

Six—map, terrain and signature data. A live feed is a dependency; the same data cached before deployment is not. The distinction is exactly the temporal one at §6.4, and it costs storage, which is cheap, rather than permission, which is not.

Seven—the relay layer. Aerial relay nodes belonging to the operator’s own force pass the test. Where the relay function is provided as a commercial service, it does not.

¹³Starlink satellite services in Ukraine’ https://en.wikipedia.org/wiki/Starlink_satellite_services_in_Ukraine accessed 5 August 2026. An encyclopaedic source is used here for uncontested chronology only; every proposition on which the argument turns is carried by a source at nn 203 to 205.

¹⁴ibid.

¹⁵PBS NewsHour, ‘Elon Musk’s refusal to provide Starlink support for Ukraine attack in Crimea raises questions for Pentagon’ (*PBS NewsHour*) <https://www.pbs.org/newshour/economy/elon-musks-refusal-to-provide-starlink-support-for-ukraine-attack-in-crimea-raises-questions-for-pentagon> accessed 5 August 2026.

¹⁶Belfer Center for Science and International Affairs, ‘Starlink and the Russia-Ukraine War: A Case of Commercial Technology and Public Purpose?’ <https://www.belfercenter.org/publication/starlink-and-russia-ukraine-war-case-commercial-technology-and-public-purpose> accessed 5 August 2026, recording the operator’s position that the network was to remain civilian, its restriction of military uses including drone control, and the subsequent establishment of a separate government product line.

¹⁷Joscha Abels, ‘Private infrastructure in geopolitical conflicts: the case of Starlink and the war in Ukraine’ (2024) 30 *European Journal of International Relations* 842.

¹⁸Rivero (n 4).

Eight—frequency allocation and deconfliction. Where operation requires an assignment from another authority, that authority can withhold it. Ukrainian experience records the related and non-adversarial failure of electronic fratricide, in which friendly emissions degrade friendly control.¹⁹

Nine—telemetry and usage reporting. A platform which reports status, position or usage to a manufacturer discloses to a party who did not need to know, over a channel that exists whether or not the mission wants it. This fails the test in a second way: it is not only a dependency but an emission, and it belongs on both axes.

Nine A—the undisclosed or denied vendor relationship. This class was not in the first version of the audit and it is a distinct failure mode rather than a variant of the others. The platforms demonstrated by the People’s Liberation Army at *Golden Dragon 2024* were identified in reporting as a Unitree B1; the manufacturer states that it does not sell its products to China’s military, and how the platforms were acquired has not been established.²⁰

Whatever the facts, the *structure* is what the audit is for, and it is worse than an acknowledged dependency rather than better. An operator whose relationship with a vendor is unacknowledged has no contract, no support obligation, no assurance of continued supply, no route to demand disclosure of what §9.1’s classes two and nine contain, and no standing to complain when any of it stops. The vendor cannot withdraw a service it does not admit providing—but it also cannot be required to keep providing one, cannot be audited, and cannot be held to a specification.

This inverts the framing the rest of this Part uses. Elsewhere the concern is a supplier who *may* withdraw. Here the supplier has, on its own account, never supplied, and the platforms are in service regardless. Grey acquisition removes the party who could withdraw and replaces them with nobody, which solves the withdrawal problem by making the support problem total. Any state acquiring capability by this route is buying the hardware and none of the sustainment, which is a strong practical argument for closure independent of anything in Part II: a platform that needs nothing from its manufacturer is the only kind worth obtaining this way.

Ten—component and spares supply. This one **passes**, and saying why is the best test of whether the criterion does real work rather than merely condemning. Components are consumed before the mission, and a commander who holds spares holds them. A supplier’s later change of mind does not reach into an engagement already under way. That the supply chain is a serious strategic vulnerability is true, and it is not what the withdrawal test measures; the test isolates services discharged during commitment, and extending it to logistics would make it useless by making everything fail.

9.1A A deployed system that is dependent by design

The audit above is applied to a composite. One real system is worth setting beside it, because it is deployed at scale, it does most of what this paper’s architecture does, and its dependency is not an oversight but its selling point.

Anduril’s autonomous surveillance towers perform persistent unattended detection, classification and tracking. More than three hundred had been deployed by September 2024, with over two hundred fielded by US Border Patrol in the preceding two years; the extended-range variant is an eighty-foot mast whose sensor head is stated to classify and track to about seven and a half miles and to detect autonomously beyond five.²¹ That is a working existence proof for persistent autonomous sensing at a scale nothing in the military

¹⁹ibid.

²⁰Army Recognition, ‘Chinese Army Demonstrates Machine Gun-Equipped Robotic Dogs in Joint Exercise with Cambodia’ <https://www.armyrecognition.com/archives/archives-land-defense/land-defense-2024/chinese-army-demonstrates-machine-gun-equipped-robotic-dogs-in-joint-exercise-with-cambodia> accessed 5 August 2026, reporting the *Golden Dragon 2024* exercise of 16 to 30 May 2024, a reconnaissance variant of about 15 kg quoted at two to four hours with real-time image transmission, and a heavier variant above 50 kg carrying up to 80 kg and a mounted assault rifle. On the manufacturer’s position and the unresolved acquisition route, Kharon, ‘At Unitree Robotics, a Star Chinese Firm, the Military Connections Keep Mounting’ <https://www.kharon.com/brief/unitree-robotics-china-pla> accessed 5 August 2026. Nothing in this paragraph depends on which account is correct; the audit point concerns the structure of an unacknowledged relationship, not its facts.

²¹ASD News, ‘Anduril Deploys 300th Autonomous Surveillance Tower (AST), Advancing Capability Border Security’ <https://www.asdnews.com/news/defense/2024/09/27/anduril-deploys-300th-autonomous-surveillance-tower-ast-advancing->

ground-robotics literature matches.

It is also, by design, the opposite of closed. The tower’s stated value is that it feeds a networked command platform, and the capability being sold is the combination rather than the sensor.²² Applying §6.1: the mission as specified fails if the network layer is withdrawn, because the operator’s product is the fused picture and not the tower. That is a deliberate and commercially sound architecture for a border-security customer who owns the network, operates in uncontested spectrum, and has no adversary able to interdict it.

Two things follow for this paper and the second is the more important.

The dependency is not a defect of engineering but a consequence of the customer. Nothing in the audit criticises the design. It establishes that the design encodes an assumption—an available, uncontested, operator-owned network—and that the assumption does not hold on contested ground. A system procured on the strength of border performance and employed against a peer inherits an architecture chosen for conditions that no longer obtain, and the withdrawal test is the tool that makes that visible before the procurement rather than after.

And the range comparison is the honest measure of what closure costs, stated carefully because the obvious version of it is wrong. The tower classifies at thousands of metres. The emplaced architecture’s *always-on cue* reaches tens of metres per node, though its imaging tier once woken reaches much further and its cue footprint extends with distributed nodes; what it cannot do is notice anything that stays outside that footprint, unless it pays for a scheduled sweep at the prices given at §16.1. The tower buys that range with height, with continuous power, and with active sensing—the three things the emplaced system gives up to remain unfound. Part VIII §28 states the boundary in those terms, and the comparison should be kept in view whenever this paper’s architecture is described as sufficient.

9.2 The count

Of the eleven classes now enumerated, **ten are discharged during the mission and one before it.** Of the ten, between four and seven are supplied by a party other than the operator, the range reflecting configuration: a platform on fibre with cached maps and on-platform inference sits at the low end; one on a commercial satellite path with delivered updates, off-platform inference and live map data sits at the high end. *[Derived: the enumeration is the author’s, from the control methods described at n 200 and the classes at §9.1. It is a count of a composite, not a measurement of a product.]*

Two observations follow, and the second is the more important.

The count is **larger than the discussion assumes**, because the discussion concerns the control link, which is one of ten. A specification that secures the control link and leaves the other eight unexamined has addressed between one tenth and one quarter of the exposure.

And the exposure is **not principally about jamming**. Eight of the ten can be lost without an adversary doing anything at all. That is the finding of this Part, and it is why closure is treated here as an architectural property rather than as a hardening measure. Hardening addresses attack. It does not address a supplier’s terms of service.

10. What the audit does not establish

Three limits, and they bound the result rather than qualifying it away.

The stack is a composite and not a product. No fielded system has been audited against its own specification, because the specifications are not public. The audit therefore establishes the shape of the

capability-border-security accessed 5 August 2026; Anduril, ‘Anduril Launches Extended Range Sentry Tower (XRST) to Advance Counter Intrusion Missions’ <https://www.anduril.com/news/anduril-launches-extended-range-sentry-tower-xrst> accessed 5 August 2026, for the eighty-foot mast and the classification and detection ranges. Manufacturer figures.

²²Anduril, ‘Sentry’ <https://www.anduril.com/sentry> accessed 5 August 2026, describing the platform’s function in combination with the vendor’s command software.

exposure and a method for measuring it. It does not establish the count for any particular platform, and a reader with access to a real specification should expect a different number.

Disclosure limits the method. Classes two and nine—entitlement checks and telemetry—are the classes least likely to appear in documentation supplied to a purchaser, and they are the two whose operational consequence is most severe. An audit conducted on documentation understates, and by an unknown amount. That is an argument for requiring the disclosure as a procurement condition, which is a policy point rather than an engineering one and belongs to the companion legal paper.

Closure is not a virtue in itself. A closed system that cannot do the job is worse than an open one that can. This Part establishes that dependency exists and is under-examined. It does not establish that the dependent capabilities are obtainable otherwise. That argument is made at Part II §7.2 and specified at Part V, and if it fails, the correct conclusion is that the dependency is worth its price rather than that it is not there.

Part IV—The Emission Budget

11. What this Part can and cannot establish

This section aims to convert emission posture from a description into a design quantity. It states, for each posture on the ladder at Part II §7.1, what is radiated and what that implies for an adversary attempting first to detect and then to locate the emitter.

The limit has to be stated before the argument rather than after it, because it governs how the figures should be read. **Absolute detection ranges cannot be given.** They depend on the intercepting receiver’s aperture, noise figure, integration time and processing, on terrain and on propagation conditions, and the relevant capabilities of a peer adversary are not public. Any paper offering a specific kilometre figure for time-to-detection against an unspecified opponent is offering a number it cannot support.

What can be given, and what is given here, are the **scaling relationships**. These follow from propagation and detection theory rather than from assumptions about a particular opponent, they are the quantities a designer can actually act on, and they are sufficient to rank the postures against one another. Where a relationship is standard it is cited. Where a consequence is the author’s arithmetic it is marked as derived.

12. The asymmetry that governs everything

The founding fact is that **interception is easier than communication**, and it is easier by a margin that is structural rather than incidental.

A communication link must deliver enough signal-to-noise ratio at the intended receiver to recover the message with an acceptable error rate. An interceptor need only establish that energy of a given character is present. The second task tolerates a far lower signal-to-noise ratio than the first, and the interceptor is free to spend aperture, integration time and processing on the problem in quantities the tactical receiver cannot carry. The consequence is that the radius within which an emission can be *detected* exceeds the radius within which it can be *used*, generally by a large factor.

This is why emission control doctrine reaches the conclusion it does. The doctrinal statement is unusually blunt for a body of official writing: the most effective means of evading signals intelligence detection is not to emit during periods of acute threat.²³ It is not a counsel of perfection but a statement about which variable dominates. The United States Army has been developing emissions control standard operating procedures for ground formations on exactly this reasoning, on the express footing that the always-connected posture of the Iraq and Afghanistan campaigns does not survive contact with a peer.²⁴ Naval practice reached the

²³The proposition is standard in emission control doctrine. For the ground-forces statement see US Army (n 7); for the naval statement see US Naval Institute (n 11).

²⁴US Army (n 7). See also ‘Convergence and Emission Control’ *Military Review* (November–December 2023) 59 <https://www.armyupress.army.mil/Portals/7/military-review/Archives/English/Nov-Dec-23/Tetreau/Convergence-and-Emission-Control-UA.pdf> accessed 5 August 2026.

position earlier and holds it harder, because a submarine’s survival has always depended on it.²⁵

Three properties of the emitter are available to a designer, and they act on different parts of the interceptor’s problem.

Duty cycle acts on opportunity. An interceptor cannot detect what is not being transmitted. Halving the fraction of time a system radiates halves the expected number of intercept opportunities in a given period, and this is the only one of the three that scales without limit: at zero duty cycle the problem disappears.

Waveform acts on the interceptor’s processing burden, and the mechanism should be named precisely because it determines what the other two levers are worth. Optimal detection of a spread-spectrum signal by a receiver that does not hold the spreading sequence is impractical, so the realistic interceptor uses a **radiometer**—an energy detector, which measures received energy in a band over an interval and declares a signal when it exceeds a threshold.²⁶ The asymmetry follows directly: the intended receiver applies the full processing gain of a matched filter, while the interceptor cannot, and is further penalised by having to accept a wide bandwidth and the noise that comes with it.²⁷

The asymmetry can be stated as a quantity rather than as an impression, and it should be, because it is what the rest of Part IV rests on. Urkowitz’s treatment of energy detection gives the structure: for a radiometer observing a band of width W for a time T , the decision statistic is central chi-square with $2TW$ degrees of freedom when no signal is present and non-central chi-square with the same degrees of freedom when one is, the non-centrality parameter being the ratio of signal energy to two-sided noise spectral density.²⁸ The consequence is the standard one and the whole of the argument: an intended receiver holding the spreading sequence applies a matched filter and gains **in proportion to TW** , while an interceptor without it must use the radiometer and gains only **in proportion to the square root of TW** .

That square root is the interceptor’s structural penalty. Doubling the time-bandwidth product doubles the intended receiver’s advantage and multiplies the interceptor’s by about 1.4. It also identifies which variable a designer should attack, because T is the interceptor’s integration time and the one term the emitter controls directly: a transmission short enough and unpredictable enough that the interceptor cannot know when to integrate collapses T towards zero, and with it the only term by which the radiometer was ever going to close the gap.

Two things follow that the paper relies on. Low probability of intercept means a **raised cost, not an impossibility**, and against an opponent with wideband capture and retrospective processing an assumption of safety is unsound. And because the radiometer works by *integrating energy over time*, transmitting in pseudo-random short bursts is the direct counter: it denies the interceptor the integration interval needed to cross its threshold.²⁹ That is the technical basis for the claim at §13 that persistence rather than power is the governing quantity, and it is why the two levers compound—a short burst is worth more when the interceptor cannot predict when to look for it.

Directivity acts on geometry. Energy radiated into a narrow beam is not available to a receiver off-axis, and the reduction in sidelobe energy is the reduction in what an interceptor outside the beam can work with. It is the least discussed of the three in the ground-robotics literature and the most useful for a static platform, for a reason specific to this mission: a system that does not move can be aligned once, at emplacement, on a bearing to its own control station, and thereafter never needs to radiate in any other direction. A mobile platform cannot make that trade.

²⁵US Naval Institute (n 11).

²⁶Don J Torrieri, *Principles of Spread-Spectrum Communication Systems* (Springer), on the impracticality of optimal detection without the spreading sequence and the consequent use of the radiometer as the practical interceptor.

²⁷US Department of Defense, *Performance Metrics for Low Probability of Intercept Communication Systems* (Defense Technical Information Center, ADA277346) <https://apps.dtic.mil/sti/tr/pdf/ADA277346.pdf> accessed 5 August 2026.

²⁸Harry Urkowitz, ‘Energy Detection of Unknown Deterministic Signals’ (1967) 55 *Proceedings of the IEEE* 523. The square-root relationship for the radiometer against the linear one for the matched filter is the standard consequence of that result and is not original here.

²⁹*ibid*, on time-hopped transmission of frequency-hopped and direct-sequence waveforms in pseudo-random short bursts reducing an interceptor’s ability to integrate sufficient energy to cross a radiometric detection threshold.

13. Locating, which is a harder problem than detecting

Detection tells an adversary that something is transmitting. It does not tell them where it is, and the gap between the two is operationally decisive for a system whose survival depends on not being found.

Geolocation by angle of arrival requires bearings from separated receivers and an intersection. Geolocation by time difference of arrival requires the same emission to be captured at several receivers with a common time reference. Both require the emitter to be transmitting for long enough, and often on enough separate occasions, for a solution to be formed and confirmed. Both are degraded by short transmissions, by frequency agility across the observation, and by multipath in broken or urban terrain.

Two design consequences follow and they are the practical content of this Part.

The relevant quantity is not power but persistence. A high-power transmission of a few tens of milliseconds presents a poorer geolocation problem than a low-power transmission sustained for minutes, because the second offers repeated independent observations and the first may offer one. Designers optimising for low transmit power while accepting a continuous link have optimised the wrong term. *[Derived: this follows from the observation requirements above; the magnitude of the advantage depends on the adversary's capture architecture and is not quantified here.]*

Predictability is a second axis of exposure and it is usually forgotten. A system transmitting in a fixed schedule offers an interceptor the ability to prepare: to point, to allocate receivers, and to capture the whole of a window it knows is coming. The same total emission distributed unpredictably is a materially harder problem. Where a scheduled window is operationally necessary, the schedule should be part of the cryptographic material rather than a published constant.

14. The postures compared

The ladder at Part II §7.1 can now be ranked on the quantities that matter, rather than on the intuition that less connection is safer.

Posture 0, no link, radiates nothing. It cannot be detected, located or jammed through the control path, because there is no control path. It surrenders recall entirely, which is the subject of Part VII and the whole of the companion legal paper.

Posture 1, receive only, also radiates nothing, and this is the point the field has under-used. A receiver is not detectable by direction-finding of an emission it does not produce. The platform can be commanded, constrained and told to stop, while presenting the same electromagnetic profile as posture 0. What it gives up is confirmation: the operator cannot know the command was received, and must design for that. Part VII §25 argues this is a far smaller sacrifice than it appears and that posture 1 is the single most valuable design choice available for the holding mission.

Posture 2, duty-cycled two-way, is where the three levers act together and where most real design work sits. Its exposure is the product of duty cycle, window predictability, waveform and directivity, and each is independently controllable. A system emplaced statically can push all four favourably at once.

Posture 3, continuous two-way, offers the interceptor everything: persistence, opportunity, and time to form and confirm a solution. It buys full interactivity, and against a competent electronic warfare effort it is the posture that the reported degradation of control in Ukraine describes.³⁰

Posture 4, tethered, radiates nothing from the link and gives full interactivity, which makes it attractive and explains its adoption. Its costs are physical rather than electromagnetic: a line that can be cut or followed, a limit on range, and a constraint on movement. For a static holding platform, whose range from its control station is fixed at emplacement and which does not intend to move, the mobility cost is close to zero—which is a general feature of this mission and is the reason its design space differs from the one the discipline has been working in.

³⁰Rivero (n 4), reporting that operators struggle to maintain reliable links under intense jamming and that electronic fratricide degrades radio control.

15A. The reporting problem, and a rule for when to break silence

Part IV has treated emission as a liability throughout. That is incomplete, and the omission matters enough to correct here rather than concede later.

A holding system observes. Over thirty days on a contested approach it accumulates the single most valuable product it generates: a record of what crossed, when, in what strength, and in what pattern. At posture 0 or 1 none of that reaches anyone until the platform is recovered, and a platform that is destroyed takes it with it. **The system’s intelligence product and its survivability are in direct competition, and every watt-second spent reporting is bought from the same account that pays for not being found.**

No treatment of this trade appears in the ground-robotics literature, which discusses links as control channels and not as reporting channels. The trade is nonetheless decidable, and it decides on a comparison the operator can actually make.

Report when the expected value of the information exceeds the expected cost of the disclosure.

Both sides are estimable before emplacement, which is what makes the rule usable rather than merely correct. The information side is a function of what the observation would change: a detection that alters a commander’s decision has value; one that confirms what is already known has almost none. The disclosure side is a function of the emission’s persistence and predictability, per §13, and of what the loss of this platform would cost—which for a cheap emplaced system is low, and which is itself an argument for cheap emplaced systems.

Three consequences follow and they are design consequences rather than observations.

Report thresholds should be set at emplacement, in the same act that sets the authority envelope. The commander who emplaces the system knows what would change their mind, and that knowledge is perishable. Encoding it as a reporting threshold—strength, direction, rate, or a named pattern—converts a judgement that would otherwise have to be made by a machine into one made by a person in advance, which is the same move Part V §16.3 makes for engagement authority and it should be made for the same reason.

Silence is data and the operator should be told what it means. A system reporting only on threshold breach transmits nothing on an ordinary day, and an operator receiving nothing learns that nothing exceeding the threshold occurred—provided the system is known to be alive. That proviso is the whole difficulty, and the answer is a periodic liveness indication whose absence is informative, which is a small, predictable emission and therefore the worst kind. The resolution is to accept the ambiguity: treat a silent platform as alive until its authority expires, and let expiry rather than confirmation bound the consequence.

A report is worth more when batched. Thirty days of threshold events transmitted in one burst on recovery of a corridor, or on a single scheduled window near the end of the emplacement, carries nearly the full intelligence value at a fraction of the exposure—because the value of most of these observations is retrospective and pattern-level rather than immediate. Only observations that would change a decision *now* justify breaking silence *now*, and the rule above is what distinguishes them.

A transmit chain is a substantial and often dominant consumer in a small platform’s electrical budget, and it consumes in proportion to how much it is used. Posture is therefore not only a survivability setting but the largest single lever on the power budget of a system whose other loads have been reduced—which is what Part VI §18 proposes and §19 prices.

Part V—The Architecture Specified

16. A holding platform, described so that it can be argued with

This section sets out the architecture the preceding Parts imply. It is specified rather than described, because a proposal which stops at description leaves the questions that decide whether it works—what draws power,

what holds the keys, what may fire and on whose authority—for somebody else to answer. Every element is stated so that a reader who disagrees can identify the element they disagree with.

The design case throughout is a single armed platform emplaced to deny an approach for thirty days, with a control station five to fifteen kilometres to the rear, operating in terrain and at a latitude comparable to eastern Ukraine. Where a design choice depends on that case, the dependence is stated.

16.1 Sensing

Tier one, always live: seismic and acoustic. A geophone and a small acoustic array, both passive, both cheap in power, and both effective against the signatures that matter—footfall, tracked and wheeled movement, and the acoustic signature of small aerial platforms. This tier is the reason the power budget in Part VI closes, and its power figures are given there rather than here.

Its **range** must be given here, because it constrains the geometry and the field results are less flattering than the concept. Field trials of seismic sensors in a wireless sensor network report reliable footstep detection to about 30 m, with vehicle detection tested at 10, 20 and 30 m.³¹ Other work extends footstep detection to about 94 m under favourable conditions, and evaluates stealthy footfall at about 51.5 m.³² The physics sets the shape of the decline: Rayleigh surface waves fall off as the inverse of range while body waves fall as its square, so a surface-coupled geophone does better than the raw energy argument suggests, and site-specific frequency-dependent soil attenuation makes any single figure a property of the ground rather than of the sensor.³³ Discrimination in noise is a solved problem to a useful standard—statistical measures of the seismic amplitude distribution such as kurtosis identify footfall against a noisy background.³⁴

Two design consequences follow and the second is the more important.

The engagement sector for personnel is tens of metres, not hundreds. That is consistent with the mission—a fighting position denies a specific approach, a treeline gap, a trench junction, a road—and it should be stated rather than glossed, because a reader who assumes a several-hundred-metre seismic sector will conclude the architecture covers more ground than it does.

The cue radius is not the system’s detection range, and confusing the two is the easiest error to make here. Tier three, once woken, can identify at far greater distance than tier one can trip—an uncooled imager with reasonable optics reaches hundreds of metres against a human. But nothing outside the seismic and acoustic footprint *wakes* it, so the system’s effective detection envelope is the union of its tier-one footprints and not the reach of its best sensor. A holding system built this way is blind, by construction, to anything that stays beyond the cue radius.

That limitation is purchasable rather than absolute, and the cascade arithmetic of Part VI §18.3 prices it. A scheduled imaging sweep covers ground the passive tiers cannot reach, at the cost of one tier-three event per sweep: *[Derived, at 10 W for 30 s per sweep.]*

Sweep cadence	Energy per day	Over 30 days	Harvesting surface	Multiple of the passive budget
Hourly	2.0 Wh	60 Wh	5.4 Wp	×2.2
Every 10 minutes	12.0 Wh	360 Wh	32 Wp	×13
Every 2 minutes	60 Wh	1.8 kWh	162 Wp	×66

³¹Gökhan Koç and Korkut Yegin, ‘Footstep and Vehicle Detection Using Seismic Sensors in Wireless Sensor Network: Field Tests’ (2013) 9 *International Journal of Distributed Sensor Networks* 120386.

³²G Succi and others, ‘Problems in seismic detection and tracking’ in *Unattended Ground Sensor Technologies and Applications II* (SPIE 2000), and the related range-limitation work on seismic footstep detection reporting detection to approximately 94 m and stealthy footfall at approximately 51.5 m.

³³ibid. Vertical and longitudinal waves diminish as the inverse square of range; Rayleigh surface waves as its inverse. Excess non-geometric attenuation is site-specific and frequency-dependent, and is determinable in situ.

³⁴Koç and Yegin (n 31).

An hourly sweep roughly doubles the budget and remains within everything Part VI claims. A ten-minute cadence costs thirteen times the passive budget, needs about 1.4 kg of cells or a 0.2 m² panel, and is still tractable—but it has forfeited the concealment argument at §19.1, because a surface that size is no longer small. Below a few minutes the architecture has become the continuously-imaging system it was designed not to be. **The right way to state the design is therefore that passive cueing is free and active search is metered**, and a commander choosing a cadence is buying detection envelope by the watt-hour.

Detection should be distributed and effect should be concentrated. If tier one reaches 30 m and the platform carrying the weapon is expensive relative to a geophone, the sensible arrangement is several cheap seismic nodes emplaced across the approach, meshed on the private link of §16.6, feeding one armed platform sited for its arcs rather than for its sensing. This separates the coverage problem from the effect problem, extends the covered frontage at low cost, and has a further property: the armed platform is no longer the thing that must be closest to the enemy. It is a direct application of the tiering principle to geography rather than to time, and it is how the architecture should be built.

Tier two, woken by tier one: magnetic and passive infrared. Confirmation and coarse classification. Woken on a tier-one trip, live for seconds, then returned to sleep.

Tier three, woken by tier two: electro-optical and thermal imaging with on-platform inference. Identification, discrimination and the material for a release decision. This is the expensive tier and the one which must not run continuously, which is the whole of the argument at Part VI §18.

Not present: active sensing. No radar and no laser ranging in the baseline. Both are emissions in the sense of Part IV and both are detectable at ranges greatly exceeding their useful ones. They are available as a commanded mode where the tactical situation justifies the exposure, and they are off by default.

16.2 Inference, on the platform

Classification and discrimination are computed locally. This follows from Part III and needs no separate justification, but two consequences of it belong here.

The model must be **fixed at emplacement**. A system with no update path cannot be improved in place, so it must be adequate when it is left, and the honest implication is that it will be worse than a supervised system at the end of thirty days than at the beginning, because the adversary adapts and the system does not. Part VII treats this as a risk rather than absorbing it.

And the model must be **small enough for the power budget**, which is a real constraint and not a formality. The tiering at §16.1 is what makes it tractable: the expensive inference runs on candidates rather than on a video stream, and the number of candidates in thirty days is a small number.

16.2A The static-frame dividend, and the displacement problem it conceals

Two consequences follow from the platform not moving. The discipline treats neither, because it designs for platforms that do.

The dividend. A sensor that does not move can subtract its own background. Whatever was in the scene at emplacement and has not changed is not a target, and establishing that requires no simultaneous localisation and mapping, no continuous ego-motion estimation, and no continuous active ranging. A moving platform must solve all three continuously, because it cannot otherwise separate its own motion from the world's; the standard answer is lidar, which supplies a metric point cloud against which motion can be resolved. A static platform obtains the same discrimination from a survey performed once.

That saving has not been counted anywhere in this paper's parent literatures and it is substantial. Lidar is an active emitter—a laser range-finder is precisely what §16.1 excludes on emission grounds—and it is a continuous load of watts rather than milliwatts. Stasis removes it twice over, once from the emission budget and once from the power budget, and the tiered arithmetic at Part VI §18.2 is only available because no continuous ranging is running. **The reason a holding platform can be powered by a surface the size of a postcard is partly that it does not have to keep asking where it is.**

The problem the dividend conceals. It is bought with an assumption—that the frame established at emplacement is still the frame the platform occupies—and the assumption can fail with nothing detecting that it has. A near-miss detonation, ground subsidence, freeze–thaw heave, a vehicle clipping the platform, or a failed recovery will displace it. Its sensing then reports a world in which everything has moved at once, which is the classic failure of a static-frame perception system, and the consequence here is worse than blindness.

It is worse because of what §16.3 rests on. The geometric authority constraint is defended there on the ground that it is a property of how the machine is bolted down rather than a parameter held in software, and therefore cannot be moved by a fault, a corruption or an intrusion. That is true and it is insufficient. **A displaced platform still enforces its sector mechanically, and the sector now covers different ground.** The constraint has not failed. It has silently relocated, and every argument in this paper that treats physical geometry as the reliable constraint is exposed to it. That is a defect in the specification as drafted and it is recorded here rather than in the Appendix because the answer belongs in the design.

The answer is relative rather than absolute, which is what makes it available at all. At emplacement the platform surveys not its own position but the *inter-feature geometry* of the static scene: the pairwise ranges and included angles among fixed features—a wall, a treeline, a culvert, a burnt-out hull. Those quantities are invariant under displacement of the observer. The absolute bearings to those features are not. The two cases are therefore distinguishable from the platform’s own sensing and nothing else:

the inter-feature pattern is preserved but every bearing has shifted → **the platform moved** the
inter-feature pattern itself has changed → **the world changed**

No external reference is required to make that distinction. No satellite positioning, no beacon, no survey party, no second platform. It is self-localisation that passes the withdrawal test of §6.1, which is not true of any of the alternatives, and that is the reason to prefer it over an inertial unit or a positioning fix even where those are available.

Three consequences, and the first is a requirement rather than a suggestion.

Displacement detection must gate weapon release. On detecting that it has moved, the platform reverts to safe and requires positive re-authorisation. It does not attempt to re-derive its sector and continue, because the ground its arcs now cover has not been surveyed by anyone and the whole justification for the geometric constraint was that a human chose the geometry. This is the correct failure direction and it is consistent with §26: ambiguity resolves towards inert.

Self-survey at power-on makes robotic emplacement possible, which is worth more than it first appears. A platform that establishes its own reference geometry when switched on does not need a human surveyor to site it, so it can be carried forward and set down by the logistics robots of §20. Part VII §26A identifies the emplacement, exchange and recovery cycle as the architecture’s real vulnerability precisely because those are the moments a human is forward. Removing the human from the emplacement is therefore not a convenience; it attacks the one thing an adversary can reliably attack.

The cost is close to nothing. Comparing a stored feature geometry against a fresh observation is a small computation performed on imagery the tier-three wake has already captured for classification. It adds no sensor, no emission and no continuous load, and it runs on the escalations the budget has already paid for.

One limit is real and bounds the method. It requires the scene to contain static features to survey. On genuinely featureless ground there is nothing to hold a pattern, and the fallback is a mechanical inclination and displacement switch, which detects gross movement and nothing subtle.

16.3 The authority envelope

What the platform may do is bounded by three constraints set before emplacement, and the choice of constraints is deliberate.

A geometric constraint. Engagement is permitted within a defined sector and range, fixed at emplacement by physical alignment and by survey rather than by a coordinate held in software. A boundary that is a

property of how the machine is bolted down cannot be moved by a fault, a corrupted parameter or an intrusion. It can, however, be moved by the machine being moved, which is the defect §16.2A identifies and answers: the constraint is only as good as the platform's knowledge that it still occupies the ground the sector was surveyed against, and detection of displacement must therefore gate release.

A temporal constraint. Authority expires. The platform reverts to an inert state at a time set before emplacement and requires positive action to resume. Expiry into inertness rather than into continued operation is the correct default, and Part VII §26 argues the point.

A behavioural constraint, not a categorical one. The release criterion is stated in terms of observed conduct within the sector rather than in terms of target identity. This is the most contestable choice in the specification and it is made for a reason that should be exposed rather than buried: identity classification by a fixed model on a thirty-day-old picture is the least reliable thing the system does, and building release authority on top of it compounds the least reliable element. Conduct within a defined geometry is a weaker criterion than identity and it is a more robust one. Whether it is a *lawful* criterion is not an engineering question, it is the central question of the companion legal paper, and this specification does not pretend to answer it.

16.4 Release interlocks

Two, and they are of different kinds by design.

Physical safing that is not under software control. A mechanical arrangement that must be removed for the weapon to function, and whose state cannot be changed by any fault, corruption or intrusion in the computing system. The reason for insisting on the physical form is that every software interlock shares a failure domain with the thing it constrains.

A second, independent authorisation channel for the transition from safe to live, distinct from the sensing and inference chain, so that no single failure both generates a candidate and authorises action on it.

16.5 Cryptographic material and the key ceremony

Keys are generated and loaded before deployment by a procedure the operator conducts, on equipment the operator holds, with no external service consulted at any point. There is no certificate authority to reach, no attestation server to satisfy, and no vendor in the trust path.

Capture is the design case. A platform emplaced in contested ground and left for thirty days may be taken, and a specification that assumes otherwise is not serious. Three measures follow: keys are unique per platform or per small group, so that a compromise does not propagate to a fleet; key material is held so that it can be erased on tamper detection, with erasure the default response to an ambiguous indication; and the mesh is designed to degrade to isolated operation on the loss of any member rather than to fail. The honest statement is that these reduce the consequence of capture and do not eliminate it, and Part VII returns to the point.

16.6 Communication

Default posture 1, receive only, per Part IV §14. The platform listens on a schedule that forms part of the cryptographic material and never transmits.

A private mesh across the emplaced group, at posture 2 when commanded. This is where the coordination in the networking argument lives: detections shared, tracks handed off, a common local picture maintained across a group of platforms the operator emplaced, on keys the operator holds. It satisfies the withdrawal test because there is no supplier in it.

Directional antennas aligned at emplacement, per Part IV §12. A static platform can make this trade and a mobile one cannot.

A fibre option where the tactical situation permits laying it, which converts the link to posture 4 and removes the electromagnetic problem at the cost of a physical trace.

16.7 Sustainment

Ukrainian practice on maintenance should be adopted rather than improved on, because it was arrived at under conditions that select hard for what works. Brigade-embedded workshops of ten to twelve engineers, diagnosing, fabricating parts, and reconfiguring in hours rather than routing to a depot.³⁵ Mobile repair capability at a cost of roughly USD 36,000 per vehicle carrying two technicians, a 3D printer and welding equipment with two days of autonomous operation.³⁶ And the design implication that follows: the platform must be repairable by the unit holding it, which means published interfaces, no parts pairing, and no authentication of components against a vendor service. A platform which requires the manufacturer's permission to accept a replacement part has failed the withdrawal test through its maintenance chain, whatever its control link does.

16.8 What the specification does not settle

Three things, stated here so that Part VI is read against them.

It has not been built. The architecture closes on published component figures. That is a much weaker statement than a bench result and it is not offered as one. The Appendix sets out the build and the measurements that would settle it.

The release criterion is not defended here. §16.3 explains why conduct-in-a-geometry is more robust than identity for a fixed model. Whether a system may lawfully engage on that basis is a legal question this paper does not answer and should not be read as answering.

The energy budget is the open term. Everything in §16 is achievable with existing components. Whether it can be *powered* for thirty days is the question the specification exists to pose, and Part VI is the answer.

Part VI—Energy

17. The discipline is measuring the wrong quantity

This section aims to establish that the endurance figures published for unmanned ground vehicles are not merely inadequate for the holding mission but are measurements of a different problem, and to propose the quantities the mission requires.

17.1 The discrepancy

Fielded platforms are rated at four to eight hours for small types and eight to sixteen hours for mid-sized logistics types.³⁷ The stated ambition for the next generation is endurance measured in days rather than hours.³⁸ Against those figures, a platform holding a fighting position for more than a month is between forty and two hundred times the quoted endurance.³⁹

A second, independent data point comes from the other side of the technology and it confirms the shape. At the *Golden Dragon 2024* exercise in Cambodia in May 2024, the People's Liberation Army Southern Theater Command demonstrated armed quadrupedal platforms alongside Cambodian forces: a reconnaissance variant of about 15 kg quoted at **two to four hours of battery life** with real-time image transmission to command posts, and a heavier variant above 50 kg carrying up to 80 kg of payload and a mounted assault rifle.⁴⁰ Two states, two morphologies, two industrial bases, and the same order of magnitude. The reconnaissance figure

³⁵Rivero (n 4).

³⁶ibid, giving approximately USD 36,000 for mobile repair vehicles developed by the Lithuanian Riflemen's Union.

³⁷Unmanned Systems UK (n 2).

³⁸ibid.

³⁹*Militarnyi* (n 1).

⁴⁰Army Recognition (n 20). These are figures reported from a state demonstration and are treated on the same footing as the Ukrainian reporting at n 10.

is also a description of the architecture this paper argues against—continuous imaging and a continuous link—quoted with the endurance that architecture produces.

The reconciliation is not that the reported case is wrong. It is that **the quoted figure is a drive-hours figure**, and holding ground is not a drive-hours problem.

The energy budget of a moving ground platform is dominated by locomotion. Every other load—sensing, computation, communication, actuation of a weapon—is small beside the work of moving a mass across broken terrain. That is why endurance is quoted the way it is: for the missions the discipline has been designing for, which are logistics, engineering and assault, the locomotion term is the budget and the rest is rounding.

For a platform emplaced to deny ground, the locomotion term is approximately zero, because not moving is the mission. Every term that was rounding is now the entire budget, and the term that was the entire budget is gone. A figure that measures the second tells you almost nothing about the first.

The current procurement record confirms this from the negative side and is worth stating because it dates the argument. Ukraine’s stated 2026 objective is that **all** frontline logistics be performed by robotic systems, with 25,000 unmanned ground vehicles contracted in the first half of the year, more than double the 2025 total, against more than nine thousand ground-robot missions in March alone.⁴¹ The platform codified under NATO cataloguing standards in the same month is specified as a 300-kilogram payload with a 50-kilometre range.⁴²

Payload and range. Both are mobility quantities, and neither says anything about what the platform can do while stationary. **The field is scaling rapidly and specifying entirely in the dimensions of the mission it is scaling for**, which is the mission this paper is not about. That is not a criticism—the logistics case is real and the figures show it being met—but it is why no measure exists for the holding case, and why §17.2 has to propose one rather than adopt one.

The single reported case of ground robots taking a position without infantry was an assault rather than a hold, and it is described in exactly those terms.⁴³ The distinction this Part rests on therefore survives the most advanced fielded practice available: the discipline has learned to move things and to take things, and holding is where its metric stops working.

17.1A Morphology, which turns out to matter and is not discussed

The two data points above differ in a way worth extracting, because it produces a design conclusion that runs against where investment is currently going.

A tracked or wheeled platform at rest draws **nothing** from its drivetrain. Chocks, brakes and gravity hold it in place, and the motors are unpowered. A legged platform at rest does not have that property for free: unless it lies down, or its joints are mechanically locked, holding a posture against gravity is an active task with a continuous power cost. Legged designs also carry a substantially higher cost of transport than tracked ones, which is why their quoted endurances are short, but that is the mobility argument and it is not the one that matters here.

For the holding mission the relevant question is not what a platform costs to move. It is **what it costs to do nothing**, and on that measure the ordering of morphologies inverts. The legged platform is attracting attention because it crosses broken ground, climbs and enters buildings—genuine advantages for assault, reconnaissance and urban work. For a system whose entire mission is to remain in one place for thirty days

⁴¹Katie Livingstone, ‘Ukraine to field 25,000 ground robots in push to replace soldiers for frontline logistics’ *Defense News* (24 April 2026) <https://www.defensenews.com/unmanned/2026/04/24/ukraine-to-field-25000-ground-robots-in-push-to-replace-soldiers-for-frontline-logistics/> accessed 6 August 2026, reporting the Defence Minister’s statement of 18 April 2026 that more than nine thousand missions were carried out using unmanned ground vehicles in March alone, and the stated objective that all frontline logistics be performed by robotic systems.

⁴²*ibid*, reporting the codification of the Bizon-L, a 300-kilogram-payload logistics platform with a 50-kilometre range, under NATO cataloguing standards.

⁴³*ibid*, reporting the President’s statement of 14 April 2026 that a robotic strike unit of the 3rd Separate Assault Brigade took a fortified position in Kharkiv Oblast using aerial drones and unmanned ground vehicles alone. The operation is described throughout as the capture of a position rather than its retention.

it is the worst available choice, and the property that makes it attractive elsewhere is irrelevant here because the platform is not going anywhere.

Two design consequences follow, and neither appears in the literature on either platform type. A holding platform should be **tracked, wheeled or emplaced without locomotion at all**—there is no requirement that a system which never moves be able to. And where a legged platform is used for the hold mission for other reasons, its posture must be **mechanically latched rather than actively maintained**, because otherwise the standing cost alone will dominate the budget at §18.2 and none of the rest of this Part’s arithmetic applies. That is a requirement, not a preference, and it is stated here because it is the kind of thing that is obvious once written and expensive to discover late.

17.2 The platform metric

The quantity a designer of a holding system needs is:

watt-hours per emplaced-day, at a stated readiness level.

Three features are load-bearing.

Per emplaced-day rather than per operating hour, because the platform is emplaced continuously and the question is what a day of continuous emplacement costs.

At a stated readiness level, because the figure is meaningless without it and this is where most of the design freedom lives. A system that detects, classifies and engages within two seconds of a trip costs far more per day than one permitted twenty seconds, and the difference is not marginal. Readiness should be stated as the interval from a tier-one trip to a release decision being available, and quoted with the figure in the manner of any other engineering specification.

Watt-hours rather than a battery size, so that harvested and delivered energy enter the same account as stored energy. A system drawing 5 Wh per emplaced-day can meet it from a cell, from a panel, from an exchanged pack or from any combination, and the metric should not presuppose which.

17.3 The formation metric

Above the platform sits the quantity a commander actually allocates against:

energy per square kilometre per day of denial.

This is the more interesting of the two, because it permits robotic and human terrain denial to be compared on a single axis for the first time. Holding ground has always been an energy problem. An infantry section denying an approach consumes rations, water, fuel for warmth, and the transport energy to deliver all three, and the reason a position is expensive to hold has always been that this consumption must be carried forward through contested ground. Robotic holding is the first form of the problem whose consumption is electrical and therefore directly measurable rather than inferred.

The comparison is not made in this paper, because doing it properly requires infantry sustainment figures at a granularity the author does not have and would have to assume. Stating the metric and declining to populate it is the honest course. What can be said is that the comparison is now *possible*, that it is the right way to argue about whether robotic denial is efficient, and that the arguments currently made about cost are made in unit prices, which measure the wrong thing in the same way drive hours do.

18. The architectural inversion

18.1 Two literatures that do not cite one another

Unattended ground sensors achieve emplacements measured in months and years. They do it on primary cells, sometimes with harvesting, and the technique is not exotic: the systems are **event-driven**. A seismic or acoustic front end idles at microwatt levels between scan events and the expensive parts of the system

are woken only when something has happened.⁴⁴ Multi-modal designs—seismic, acoustic, magnetic, infrared and imaging in tiers—are the ordinary form of these systems, and the tiering exists precisely to keep the imaging tier asleep.⁴⁵

Unmanned ground vehicles are designed as vehicles. They are architected around a drivetrain, a continuously available sensing and control loop adequate for motion, and a link adequate for supervision. The power architecture follows from the assumption that the platform is doing something continuously.

For the holding mission the second architecture is wrong and the first is right, and the two fields do not appear to cite one another. **The proposal of this Part is that a holding platform should be architected as an unattended ground sensor that has effectors, rather than as a vehicle that has stopped.**

18.1A The inversion has already been fielded, and the comparison is measured rather than derived

The argument to this point has been a design calculation. It does not have to be. A fielded programme supplies the same comparison as a measurement, in the same mission role, and it should be set out before the arithmetic because it is stronger evidence than the arithmetic.

The Remotely Monitored Battlefield Sensor System and its improved variant are passive unattended ground sensors which, once emplaced, operate unattended for **up to thirty days**, on a stated minimum battery life of thirty days at **1,000 detected events per day**.⁴⁶ The comparison the programme literature itself draws is with the AN/PPS-5D ground surveillance radar, which is quoted at **approximately eight hours**.⁴⁷

Set those two side by side. Same task—detecting movement across ground. Same era of electronics. **Thirty days against eight hours is a factor of about ninety, and the difference is not battery chemistry but architecture:** one is a passive, event-driven sensor which spends nothing until something happens, and the other is an active emitter which spends continuously whether anything happens or not.

That is this paper’s central claim, measured on fielded hardware, by the organisation that fielded it, decades ago. The inversion at §18.1 is not a proposal for a new technique. **It is a proposal to apply to armed ground platforms an architecture whose advantage the unattended-sensor community measured and has been exploiting since before the ground-robotics field existed.**

Three further points follow and the third is a caution against overreading.

The 1,000-events-per-day figure is a useful calibration and it cuts against the assumption at §18.2. This paper’s budget assumes ten tier-three escalations a day. REMBASS is specified to survive thirty days while *detecting* a thousand events a day—two orders of magnitude more activity at the cheap tier. That is consistent with the cascade at §18.3, because a detection at tier one is not an escalation to tier three, and it suggests the tier-one trip rate this paper assumed at 200 a day may be low rather than high. Since tier one costs 12 mWh a day, an error of that size in that term is immaterial, which is itself a demonstration of where the budget’s sensitivity actually lies.

The direction of travel in the research is further still. The DARPA near-zero-power programme is directed at extending unattended sensor life from weeks to years, on wake-on-event front ends of exactly the

⁴⁴The event-driven architecture is the defining feature of the class. See ‘Unattended ground sensor’ https://en.wikipedia.org/wiki/Unattended_ground_sensor accessed 5 August 2026 for the general description, and for the power point specifically, seismic ground-coupled geophones operating at microwatt levels between scan events.

⁴⁵Multi-modal tiering—seismic, acoustic, magnetic, infrared and imaging—is the ordinary configuration. See GlobalSecurity, ‘Unattended Ground Sensors (UGS)’ <https://www.globalsecurity.org/military/systems/ground/ugs.htm> accessed 5 August 2026.

⁴⁶On REMBASS and Improved REMBASS as passive sensors operating unattended for up to thirty days, with a stated minimum battery life of thirty days at 1,000 detected events per day, see Federation of American Scientists, ‘Remote Battlefield Sensor System (REMBASS) / Improved Remote Battlefield Sensor System (I-REMBASS)’ <https://man.fas.org/dod-101/sy/s/land/rembass.htm> accessed 6 August 2026; and GlobalSecurity, ‘Remote Battlefield Sensor System (REMBASS)’ <https://www.globalsecurity.org/military/systems/ground/rembass.htm> accessed 6 August 2026. The sensors are powered by BA-5589/U lithium cells and the repeaters by BA-5590/U.

⁴⁷The comparison with the AN/PPS-5D ground surveillance radar, quoted at approximately eight hours against REMBASS II’s approximately thirty days, is drawn in the programme literature at n 503B.

kind §18.1 describes.⁴⁸ If that matures, the floor assumed here is pessimistic.

And the comparison is not a like-for-like capability claim. REMBASS detects and reports; it does not classify, decide or engage, and it carries none of the imaging or inference load that dominates the budget at §18.2. The ninety-fold figure establishes the architectural point and does not establish that an armed platform inherits it unchanged. What it establishes is that the passive event-driven architecture delivers order-of-magnitude endurance advantages in the field and not merely on paper, which is the premise the rest of this Part builds on.

18.2 The budget, stated so it can be checked

What follows is the author’s arithmetic on published component classes, for the architecture at Part V §16.1. Every input is stated. It is a design calculation and not a measurement, and Part VI §22 and the Appendix say what would have to be built to test it.

[All figures in this subsection are derived. Component power classes are order-of-magnitude figures for current parts; event rates are assumptions and are varied below.]

Tier one, continuous. A seismic front end with wake-on-event, a low-power acoustic detector, and a microcontroller in deep sleep with a real-time clock. Taking 500 μ W as a design point for the tier as a whole:

$$500 \mu\text{W} \times 24 \text{ h} = \mathbf{12 \text{ mWh per day}}$$

Tier two, woken. Magnetometer and passive infrared, at roughly 50 mW when live. Assuming thirty trips a day, each holding tier two live for ten seconds:

$$50 \text{ mW} \times 300 \text{ s} = \mathbf{4.2 \text{ mWh per day}}$$

Tier three, woken. Thermal imager and an edge inference processor. Published figures put an uncooled long-wave infrared microbolometer core at 2 to 5 W, with ultra-compact modules below 1 W, and a current edge inference module at 8 to 12 W under inference with peaks to 25 W.⁴⁹ Taken together the pair sits between about 3 W for a low-power core with a modest processor and about 20 W for a capable one. **10 W is used here and it is towards the optimistic end of that range;** the sensitivity is given below. Assuming ten candidates a day, each holding the tier live for thirty seconds:

$$10 \text{ W} \times 300 \text{ s} = \mathbf{0.83 \text{ Wh per day}}$$

Communication at posture 1. A duty-cycled receiver listening one second in sixty, at 30 mW while listening:

$$30 \text{ mW} \times 1,440 \text{ s} = \mathbf{12 \text{ mWh per day}}$$

This line is superseded at Part VII §25.1, which tiers the receiver on the same principle as the sensing chain and brings it to approximately 0.9 mWh per day. The flat figure is retained here because it is the conservative one and because the total moves by barely one per cent either way—tier three dominates whatever the receiver does, which is the section’s point.

Housekeeping. Power conditioning, monitoring and losses, taken at 2 mW continuous:

$$2 \text{ mW} \times 24 \text{ h} = \mathbf{48 \text{ mWh per day}}$$

⁴⁸On the near-zero-power research directed at extending unattended sensor life from weeks to years using wake-on-event front ends, see the summary of the DARPA N-ZERO programme at <https://thingslog.com/blog/2026/04/18/case-study-darpa-n-zero-near-zero-power-battlefield-sensors> accessed 6 August 2026. This is a secondary account of a research programme and is cited for the direction of the work rather than for any specific figure.

⁴⁹On uncooled long-wave infrared cores, LightPath Technologies, ‘Uncooled vs Cooled LWIR Cameras: A Guide’ <https://www.lightpath.com/insights/uncooled-vs-cooled-lwir-cameras-a-guide> accessed 5 August 2026, giving 2 to 5 W for uncooled sensors against tens of watts for cooled systems. On edge inference, ‘Jetson Orin Nano Power Consumption: Idle, 15W, MAXN Benchmarks’ (2026) <https://edgeaistack.ai/blog/jetson-orin-nano-power-consumption/> accessed 5 August 2026, giving approximately 5 W idle, 8 to 12 W under inference and up to 25 W at peak. Manufacturer and vendor figures are used here because no independent measurement of this combination in this application exists; the Appendix treats that as a gap rather than as a settled input.

Effect. Traverse, elevation and firing. A servo drawing 20 W for five seconds per engagement, at five engagements across the whole emplacement:

$$20 \text{ W} \times 25 \text{ s} = 0.14 \text{ Wh for the entire thirty days}$$

Total: approximately 0.91 Wh per emplaced-day, of which tier three is 92 per cent. Rounded, **1 Wh per emplaced-day** and **27 Wh over thirty days**.

Three things about that number deserve attention.

Effect is not the driver. The weapon costs 0.14 Wh across the whole emplacement, against roughly 28 Wh for sensing. The intuition that an armed system is an energy-hungry system is wrong for this mission; the gun is free and the eyes are expensive.

The tier-three power figure matters less than it looks. Running the budget across the published range rather than at the chosen point: *[Derived.]*

Tier-three power	Total per emplaced-day	Over 30 days	At 100 events/day
3 W	0.33 Wh	9.8 Wh	2.6 Wh/day
10 W	0.91 Wh	27 Wh	8.4 Wh/day
20 W	1.74 Wh	52 Wh	16.7 Wh/day

At the pessimistic end the thirty-day requirement is 52 Wh, or about 210 g of cells, and the December harvesting requirement rises from 2.5 Wp to 4.7 Wp—from roughly 0.015 m² to 0.028 m². Neither is a difficulty. The architecture is insensitive to a factor of seven in the largest line of its own budget, which is the most reassuring property it has, and it is a consequence of the duty cycle rather than of any component choice.

Tier three dominates, so the event rate is the sensitive parameter. At a hundred tier-three candidates a day rather than ten, the total rises to 8.4 Wh per day and 252 Wh over thirty days. At a thousand it is 83 Wh per day and 2.5 kWh, and the architecture is in trouble. The design therefore stands or falls on the **discrimination performance of tiers one and two**, which is exactly where the engineering effort should go and is not where it currently goes. That is the sharpest practical recommendation in this paper.

The comparison is the point. A conventionally architected platform running thermal imaging and inference continuously at 10 W, with a continuous radio link at 5 W, draws:

$$15 \text{ W} \times 24 \text{ h} = \mathbf{360 \text{ Wh per day}}, \text{ and } \mathbf{10.8 \text{ kWh over thirty days}}$$

The ratio at ten candidates a day is **396:1**. At a hundred it is 43:1. The claim in Part I that the inversion changes the achievable emplacement by orders of magnitude rather than by a factor is therefore supported at the assumed event rates and survives an event rate an order of magnitude worse.

18.3 Where the event rate actually comes from

The budget above assumes ten tier-three escalations a day, and the Appendix concedes that the assumption rests on plausibility rather than evidence. It can be improved on. The escalation rate is not a free parameter: it is the output of a detection cascade whose terms are separately measurable, and writing it out identifies which term the design should attack.

Let the sector receive **T** genuine events of interest per day. Let tier one produce **F** false trips per day from wind, precipitation, animals, distant traffic and sensor noise. Let **P** and **P** be the probabilities that tiers one and two respectively register a genuine event, and let **q** be the probability that tier two passes a tier-one false trip rather than rejecting it. Then the number of escalations reaching the imaging tier is

$$\mathbf{N = T \cdot P \cdot P + F \cdot q}$$

and the daily energy is the baseline plus N multiplied by the energy of one escalation. *[Derived. The cascade is elementary and the detection-theoretic framing is standard; what is offered is its application to the energy budget.]*

The paper’s assumed ten decomposes readily. Two genuine crossings a day at $P = P = 0.95$ contribute 1.8. The remaining 8.2 is false-alarm leakage: two hundred tier-one trips a day at $q = 0.041$, or a hundred at $q = 0.082$. Those are ordinary numbers for an unattended sensor in real terrain, which is the point—the assumption is recoverable from plausible parameters rather than merely asserted.

Three consequences follow, and the first is the one worth having.

The energy budget of a holding system is set by its false-alarm rate, not by enemy activity. Wherever F greatly exceeds T —which is the normal condition for an unattended sensor, since the environment is busy and the enemy is not—the second term dominates and the first is a rounding error. A sector across which a company passes daily costs almost exactly what an empty one costs. A sector in high wind, heavy rain or near a road costs several times more. That inverts the intuition a designer brings from every other part of military engineering, where cost tracks enemy activity, and it means the operating cost of terrain denial is a property of the weather and the ground rather than of the war.

Tier-two rejection is the highest-leverage parameter in the system. Since N is dominated by $F \cdot q$, halving q nearly halves the whole budget. Nothing else available to a designer has that effect: improving tier-three efficiency by a factor of two halves 91 per cent of the budget, but improving q by a factor of two halves nearly the same share at a fraction of the engineering cost, because a magnetometer and a passive infrared channel are cheap and a better imager and processor are not.

Detection range and energy cost are coupled, and they pull against each other. Raising tier-one sensitivity to extend the 30 m detection radius of §16.1 raises F in the same motion, because the additional sensitivity admits more of the environment along with more of the enemy. There is therefore an optimum sensitivity for a given energy budget, and it is below the maximum the hardware permits. That is a design result which does not appear in either parent literature: the unattended-sensor field optimises detection subject to power, and the ground-robotics field does not treat detection as an energy question at all.

The revision this forces on the Appendix should be stated. The experiment at A.1 should not measure “escalations per day”. It should measure **T, F and q separately**, because those are the three quantities that transfer to a different site, a different season and a different threshold setting, and the composite does not.

18.4 What this means in mass

At roughly 250 Wh/kg for current lithium-ion cells, the conventional 10.8 kWh is about 43 kg of cells before packaging—a substantial fraction of a mid-sized platform’s payload, which is why nobody emplaces one for thirty days.

The inverted architecture’s 27 Wh is about 110 g. *[Derived from the figures above.]* Even at the hundred-candidate rate, 252 Wh is about 1 kg. The energy problem for a holding system, on this architecture, is not a hard problem. It has been a hard problem because the wrong architecture was being powered.

19. The interlock

19.1 Why harvesting is normally dismissed, and why the dismissal fails

Photovoltaic harvesting is generally rejected for this application on two grounds: yield at high latitude in winter is poor, and a panel is a signature. Both objections are sound against a vehicle-class load. Neither survives the inversion.

Annual photovoltaic yield in Ukraine is approximately 1,100 to 1,200 kWh per kWp installed.⁵⁰ Winter yield at 50°N is a small fraction of the annual figure. *[Derived, and the assumption is stated: taking December at*

⁵⁰PVKnowHow, ‘Ukraine Solar Panel Manufacturing: Market Insights Report’ <https://www.pvknowhow.com/solar-report/ukraine/> accessed 5 August 2026, giving approximately 1,100 to 1,200 kWh per kWp per year.

approximately 2 per cent of annual yield, which is the conventional figure for a fixed optimally-tilted array at this latitude, gives roughly 23 kWh per kWp in December, or about 0.75 kWh per kWp per day. Applying a further 50 per cent derating for non-optimal tilt, soiling, snow, concealment and partial shading gives approximately 0.37 kWh per kWp per day, or 370 mWh per Wp per day.]

Against a 360 Wh daily load, meeting the budget in a Ukrainian December would require approximately 970 Wp—some five to six square metres of panel. That is the calculation behind the dismissal, and on those numbers the dismissal is right.

Against a 0.91 Wh daily load it requires approximately **2.5 Wp**, which at typical module power density is of the order of 0.015 m²—a surface smaller than a sheet of A5 paper. *[Derived from the 370 mWh/Wp/day figure above.]* At the hundred-candidate event rate and 8.4 Wh per day it requires about 23 Wp, or roughly 0.13 m². Both are concealable, terrain-matchable, and can be distributed rather than presented as a single specular surface.

The dismissal was correct about the panel and wrong about the load. This is the result referred to in Part I as the interlock, and it should be stated as a chain because each link is a consequence of the one before it. Closure removes the external service; removing the service permits posture 1; posture 1 and the tiered architecture collapse the duty cycle; the collapsed duty cycle reduces the load by between one and three orders of magnitude; and the reduced load brings the harvesting surface down to a size which is neither a signature nor a logistical burden. Harvesting then removes the resupply dependency, which is itself a form of closure. The two axes and the energy budget are not three problems. They are one problem, and solving it in one place solves it in the others.

19.2 The cold, which is the real objection

The serious objection to winter operation is not yield. It is that lithium-ion cells cannot be charged below approximately 0°C without lithium plating and permanent damage, and that their available capacity falls substantially at low temperature. A system that harvests in a Ukrainian winter must either heat its cells, which costs far more energy than the sensing does and would dominate the budget, or not charge them.

At the inverted power level a third option becomes available, and the one unattended ground sensors already use. **Do not use a rechargeable chemistry for the baseline load.**

Bobbin-construction lithium thionyl chloride primary cells carry the highest energy density of any primary chemistry—published figures reach 710 Wh/kg and 1,420 Wh/L—hold a flat 3.6 V through most of their discharge, self-discharge at one per cent a year or less at 20 °C, operate from −55 °C to +125 °C, and are quoted at operating lives up to forty years.⁵¹ They are the standard part for long-life unattended sensing for exactly these reasons.

The arithmetic that follows is almost absurd and should be stated plainly. The baseline load that must never fail—tiers one and two, housekeeping and the receiver—is 76 mWh per day, or **2.3 Wh across thirty days**. At 710 Wh/kg that is 3.2 g of cell material. *[Derived.]* In practice one would fit a standard AA-size cell of roughly 2.4 Ah at 3.6 V, about 8.6 Wh, which covers the entire thirty-day baseline nearly four times over and would run it for something over three months. The guaranteed layer of a thirty-day armed emplacement is a single AA-sized primary cell, and it works at any temperature the platform will encounter because it is never charged.

One caveat is real and it drives the rest of the design. Bobbin cells are optimised for low continuous current and cannot deliver high pulses; the standard answer in this application is a hybrid layer capacitor or supercapacitor across the cell, which supplies the burst and is replenished slowly.⁵² That is not a complication

⁵¹Tadiran Batteries, ‘Long-Life Lithium Batteries’ <https://tadiranbat.com/technical-data/long-life-lithium-batteries/> accessed 5 August 2026, giving up to 710 Wh/kg and 1,420 Wh/L for bobbin-type lithium thionyl chloride, self-discharge of one per cent a year or less, an operating range of −55 °C to +125 °C and operating lives to forty years. These are manufacturers’ published figures for the chemistry rather than an independent measurement of a particular product, and the Appendix treats the distinction.

⁵²On the low-current characteristic of bobbin construction and the standard hybrid capacitor arrangement used to serve pulsed loads, see Pkcell, ‘Comprehensive Guide to Lithium Thionyl Chloride Batteries: Operating Principles, Passivation, and IoT Hybrid Power Solution’ <https://www.pkcell.com/comprehensive-guide-to-lithium-thionyl-chloride-lisoclbatteries-operating->

so much as a confirmation, because it is the same tiering the architecture already uses: a steady microwatt source behind a store that serves the transient.

The architecture that follows is a split one, and it falls out of the arithmetic rather than being imposed on it. **Primary cells carry the guaranteed baseline**—tiers one and two, housekeeping, and the receiver—which is the part that must never fail and which totals 76 mWh per day, or 2.3 Wh across thirty days. **A small rechargeable buffer with harvesting carries the discretionary load**—tier three and effect—which is the part that is bursty, which tolerates deferral, and which can simply wait for light. If the buffer is cold and cannot charge, the system does not die. It degrades to a longer response interval, which is a readiness reduction rather than a failure, and it recovers when conditions allow.

That is a more robust design than a single rechargeable bank and it is a direct consequence of the inversion. It is also, so far as the author can establish, not how any fielded armed ground platform is built.

20. Delivery rather than generation

For loads that harvesting cannot meet—a higher readiness level, a higher event rate, a platform carrying sensors this paper has excluded, or simply a commander unwilling to accept a weather-dependent budget—energy must be brought forward. The proposal is that it be **exchanged rather than transferred**, and that the exchange be performed by the fleet that is already there.

Commercial robotics has been solving a version of this problem at scale for a decade, and its answer is instructive precisely because it does not transfer. The industry’s decision rule is settled: **swappable packs where rotations are short and staff are on site; autonomous docking where the fleet exceeds roughly fifty machines or operates continuously in unpredictable conditions**, with outdoor and large-scale deployments favouring docking and warehouse logistics favouring swap.⁵³ Many operators run a hybrid—docking for routine cycles, swap packs held for peaks and failures. The force-structure consequence is real and is worth carrying across: automated charging is credited with reducing required fleet size by up to 15 per cent, which is the commercial equivalent of the formation metric proposed at §17.3.⁵⁴

Neither branch of that rule fits the holding case, and seeing why identifies what is actually being proposed. The commercial dichotomy is between *the robot going to the energy*, which is docking, and *a human bringing energy to the robot*, which is swap. A holding platform can do neither: it cannot move, because not moving is the mission, and there is no human, because keeping humans out of the sector is the reason it is there. **The military hold case is a third category the commercial literature does not contain—the robot cannot move and the human cannot come, so another robot must bring the energy.** That is why the existing military work does not fit either: it is aerial-centric, concerned with small aircraft drawing wireless power while hovering over a ground vehicle, rendezvous points coordinated in time and space, and charging ports designed for drone swarms.⁵⁵

The commercial hybrid also inverts. There, docking is routine and swap is the exception. Here, delivered exchange is the routine fallback and the continuous source is ambient harvesting under §19—the dock, in effect, comes to the platform as daylight. Three further respects favour exchange over charging in this application, and they are specific to the environment rather than to the machine.

Time. Exchange takes minutes; charging takes hours. In contested ground the exposure of the delivering platform scales with dwell, and a logistics robot that must remain stationary beside an emplaced system for three hours is a target for three hours.

principles-passivation-and-iot-hybrid-power-solution/ accessed 5 August 2026.

⁵³Phihong, ‘Swappable Battery vs. Autonomous Docking: What Leading OEMs Are Choosing for Next-Gen Mobile Robots’ <https://www.phihong.com/swappable-battery-vs-autonomous-docking-what-leading-oems-are-choosing-for-next-gen-mobile-robots/> accessed 5 August 2026. For the underlying battery-technology position see also Yan and others, ‘A review of Li-ion batteries for autonomous mobile robots: Perspectives and outlook for the future’ (2022) 554 *Journal of Power Sources* 232324.

⁵⁴Delta-Q Technologies, ‘Building a Lights-Out Warehouse with Autonomous Charging Systems’ <https://delta-q.com/industry-news/building-a-lights-out-warehouse-with-autonomous-charging-systems/> accessed 5 August 2026, on fleet-size reduction of up to 15 per cent and charging-point costs of USD 10,000 to 50,000.

⁵⁵The Robot Report, ‘Army readies charging port for swarms of autonomous drones’ <https://www.therobotreport.com/army-readies-charging-port-autonomous-drones-swarms/> accessed 5 August 2026.

Emission and movement. Exchange requires the holding platform neither to move nor to transmit. It is the delivering platform that moves, and it is doing so anyway.

Separation of concerns. Charging couples generation to delivery: energy must be produced where the platform is. Exchange decouples them completely. Packs can be charged in a rear area from any source available—grid, generator, solar farm, vehicle alternator—and delivered forward as cargo. The contested part of the problem becomes a transport problem, and transport is the thing this fleet already does.

The fleet exists. Ukrainian ground robots conduct approximately 80 per cent of logistics in the 3rd Assault Brigade and approximately 90 per cent in Pokrovsk and Myrnyhorod, carrying loads to about 200 kg per mission, across nearly 24,500 missions in the first quarter of 2026 alone.⁵⁶ A charged battery pack is cargo of exactly the kind these platforms already carry. The proposal requires no new vehicle, no new capability and no new doctrine—only that energy be added to the list of things the resupply robot brings, and that the holding platform be designed with an exchange interface a robot can operate.

The design implications are modest and specific: a pack format standardised across the formation rather than the vendor; a mechanical interface operable by a simple manipulator without fine alignment; a connector tolerant of mud, ice and repeated blind mating; and a holding platform that continues to operate on its primary cells during the exchange, so that the exchange is never a moment of vulnerability.

21. Scavenging

This section is speculative and is labelled so throughout. It identifies a source, states an arithmetic, and sets out four engineering problems that stand between the two. It does not claim a system.

21.1 The observation

Contested ground contains a great deal of stored electrical energy. Destroyed and abandoned vehicles carry starter and auxiliary batteries. Damaged equipment carries cells. And above all, the ground is accumulating expended first-person-view aircraft, each of which carries a lithium-polymer pack.

The production figures are extraordinary and they are the point. Ukrainian production rose from a few thousand units a year in 2022 to approximately three million in 2025, with capacity in 2026 stated at more than eight million a year and monthly output exceeding 150,000 units across all types.⁵⁷ Russia is reported to be planning production of more than seven million units in 2026, some three million more than the preceding year, with estimated attack-drone output of 80,000 to 120,000 a month.⁵⁸ Ukraine's 2025 procurement alone was 4.5 million units.⁵⁹

[Derived, with the assumptions stated. Taking a combined expenditure of ten million airframes a year; a typical pack of 4S to 6S at 3,000 to 6,000 mAh, giving roughly 45 to 130 Wh and taking 80 Wh as a mid-point; and a recoverable fraction—landing intact, reachable, and with the pack undamaged—of 10 per cent. The result is one million packs and approximately 80 MWh a year, distributed across the contact line.]

Against the budget at \$18.2, a single recovered pack at 80 Wh is between two and three months of emplacement for one platform.

Two features of this source are unusual and both favour it. It is **co-located** with the mission: the packs fall on precisely the ground a holding system occupies, because that is the ground being fought over. And

⁵⁶Rivero (n 4). The mission count for the first quarter of 2026 is from GlobalSecurity, 'Over 9,000 frontline missions in March: the Defence Forces continue to expand the use of unmanned ground vehicles' <https://www.globalsecurity.org/wmd/library/news/ukraine/2026/04/ukraine-260407-ukraine-mod01.htm> accessed 5 August 2026, reporting Ukrainian Ministry of Defence figures. These are figures issued by a party to the conflict about its own performance.

⁵⁷National Security and Defence Council of Ukraine, 'Results of Ukraine's Defense Industry in 2025: FPV Drones' <https://www.rnbo.gov.ua/en/Dialnist/7370.html> accessed 5 August 2026.

⁵⁸Russia sharply increases drone production: plans to produce more than 7 million FPVs in 2026' (*Ukrainian News*, citing Bloomberg) <https://ukranews.com/en/news/1155517-russia-sharply-increases-drone-production-plans-to-produce-more-than-7-million-fpvs-in-2026> accessed 5 August 2026.

⁵⁹Ukraine to buy 4.5 million FPV drones in 2025' <https://www.yahoo.com/news/ukraine-buy-4-5-million-121037633.html> accessed 5 August 2026.

it is **counter-cyclical to every other logistics curve**: supply increases as the fighting intensifies, which is when resupply through contested ground is hardest and when a holding system most needs to be self-sufficient.

21.2 The four problems

Heterogeneity. Recovered packs vary in chemistry, cell count, capacity, connector and state. A scavenging system must accept that variation without configuration, which means a tolerant input stage capable of identifying series count and chemistry from measurement rather than from a label.

State of health on an unknown pack. A cell that has been discharged hard, crashed, and left in the cold may be usable, degraded, or dangerous, and the three are not easy to distinguish quickly. This is the hardest of the four and it is a live research area rather than a solved one. The design must fail safe: a pack that cannot be assessed with confidence is rejected rather than accepted, and the cost of over-rejection is acceptable because the supply is large.

Thermal runaway. Damaged lithium-polymer cells are a fire risk, and a system that brings unknown damaged cells inside its own structure has imported that risk. Physical separation of the scavenged bay from the platform's own energy and computing, with venting and a jettison capability, is the minimum, and it costs mass.

Mechanical interface. Recovery, connection and handling by a machine, on unknown hardware, in mud and snow. This is a manipulation problem of the kind that is unglamorous and usually decisive.

21.3 What would settle it

The proposal is testable and the test is cheap by the standards of this field. Recover a representative sample of expended airframes from a real environment; measure the fraction with packs that are physically intact; measure the fraction of those that are electrically serviceable; and characterise the assessment problem on that sample. The recoverable fraction assumed above is 10 per cent and it is an assumption, not a finding. If the true figure is 1 per cent, the arithmetic still yields 8 MWh a year across the contact line and the proposal survives. If it is 0.1 per cent, it does not.

21A. What energy buys: the discrimination exchange rate

The preceding sections treat energy as a constraint to be met. This section argues that it is better understood as a currency, that what it purchases is discrimination, and that stating the exchange rate makes a class of comparison possible which the field currently makes badly or not at all.

21A.1 The comparator nobody uses

The correct comparator for an emplaced holding system is not another ground robot. It is a **mine**.

A minefield is a holding system. It denies ground, it operates unattended for months or years, it requires no supervision, it emits nothing, it has no supply chain once laid, and it costs almost nothing per unit. On every metric this paper has proposed it outperforms the architecture at Part V, and it does so by a wide margin. Its energy consumption per emplaced-day is zero.

It also cannot tell a soldier from a child, cannot be recalled, cannot be told that the war has ended, and cannot distinguish the approach it was laid to deny from a farmer returning to a field twenty years later. That is why an entire body of law exists to constrain it, and why one class of it is prohibited outright.⁶⁰

That body of law has already answered, for the zero-energy case, two of the questions this paper has been asking of the energy-spending one, and it answered them with numbers. Amended Protocol II requires

⁶⁰Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons, Protocol on Prohibitions or Restrictions on the Use of Mines, Booby-Traps and Other Devices as amended on 3 May 1996 (Amended Protocol II), in force 3 December 1998; and Convention on the Prohibition of the Use, Stockpiling, Production and Transfer of Anti-Personnel Mines and on their Destruction (Ottawa, 18 September 1997), which prohibits anti-personnel mines outright.

remotely delivered mines to self-destruct within **thirty days** of emplacement with 90 per cent reliability, and to carry a back-up self-deactivation feature such that no more than one in a thousand remains functional after 120 days.⁶¹ It further requires mines produced after 1 January 1997 to contain at least eight grammes of iron or equivalent, so that they can be found by ordinary detection equipment.⁶²

The thirty-day figure deserves attention. This paper's design case is a thirty-day emplacement, and it was chosen from the reported Ukrainian case rather than from the treaty; the two arrived at the same number independently. That is not a coincidence so much as a convergence on the same underlying judgement about how long an emplaced effect may be left before the picture that justified it has gone stale. The law reached it in 1996 for a device with no sensing at all.

Set the two side by side and the relationship is exact. **The mine spends no energy and exercises no judgement. The architecture at Part V spends 0.91 Wh per emplaced-day and exercises some. The energy is what the judgement costs.**

21A.2 The exchange rate, stated

The budget at §18.2 can be re-read as a price list, and re-reading it that way is more informative than reading it as a constraint. *[Derived from §18.2 by reallocation; no new inputs.]*

Capability purchased	Cost per emplaced-day	Share
Presence detection only — knowing something is there	12 mWh	1.3%
Coarse discrimination — personnel against vehicle, direction, count	4.2 mWh	0.5%
Identification and release judgement	830 mWh	91.1%
Recallability — the receive channel	12 mWh (0.9 mWh tiered, §25.1)	1.3% (0.1%)
Housekeeping	48 mWh	5.3%
Effect — the weapon itself	4.6 mWh	0.5%

Four readings of that table are worth stating and none of them is available from the budget presented as a constraint.

Discrimination is 91 per cent of the price. Everything else—detecting, deciding, shooting, being recallable, staying alive—is 9 per cent between them. A system that holds ground without discriminating costs 81 mWh per emplaced-day. A system that holds ground *and looks first* costs eleven times that. Those figures use the flat receiver, and the table is drawn that way throughout for consistency with §18.2; tiering the receiver per §25.1 lowers the indiscriminating system to 70 mWh and raises the multiple to thirteen, so the figures given are the conservative ones and the argument is stronger than they state. The entire energy problem in this class of system is the cost of looking before acting.

Recallability is 1.3 per cent, and one part in a thousand once the receiver is tiered. This is the finding at Part VII §25 arriving from the opposite direction, and the agreement of the two is not a coincidence: both follow from a receiver being cheap and a transmitter being expensive. The design choice that most reduces the moral and legal hazard of the system is the cheapest line in its budget, and §25.1 shows that the emission it entails can be placed on the control station rather than on the emplaced platform. A designer who omits recall is not economising on anything.

⁶¹Amended Protocol II (n 60), Technical Annex, on self-destruction within thirty days at 90 per cent reliability and back-up self-deactivation leaving no more than one in a thousand functional after 120 days.

⁶²ibid, Technical Annex, on detectability: not less than eight grammes of iron or equivalent in a single coherent mass, for anti-personnel mines produced after 1 January 1997.

The weapon is 0.5 per cent. Restating §18.2 in this form makes the point unmissable: arming the platform is essentially free, and the constraint that keeps armed emplaced systems rare is therefore not an energy constraint. It is a policy, doctrinal and legal constraint, and any argument that treats power as the reason these systems are not yet common has misattributed the cause.

A mine is the zero-energy limit of the same design. Reduce the budget to zero and the architecture converges on a device that acts on presence alone. This is not an analogy; it is what the table shows at the bottom of its range. The two are the same class of object separated by an energy expenditure, and the expenditure is what buys the capacity to distinguish.

That last reading has a consequence the engineering can state and the law will have to decide. If the two objects differ by a quantity of energy, then the three obligations Amended Protocol II already imposes on the zero-energy case are all purchasable in the energy-spending one, and their prices are on the table above. Self-deactivation costs approximately nothing, because it is a timer. Detectability costs approximately nothing, because eight grammes of iron is already present in any such platform many times over. Discrimination—the obligation the law does *not* impose on mines, because it cannot be met by a device with no sensing—costs 830 mWh per emplaced-day. **The reason mines are constrained by marking, mapping, clearance and self-destruction rather than by a requirement to discriminate is that discrimination was not available at any price. In this class of system it is available, and the price is under a watt-hour a day.** Whether that changes what the law should require is not an engineering question; that the question is now open, and has a number attached, is an engineering finding.

21A.3 Why this reframing matters beyond this paper

Three consequences, and the third is the reason this section exists.

It supplies a **defensible efficiency argument**. The question “is robotic terrain denial worth it” has been argued in unit prices, which compare a platform to a platform and answer nothing. The question is properly whether the discrimination purchased is worth the energy and the risk, and that question has an arithmetic.

It identifies where **engineering effort should go**, and it is not where it currently goes. If 91 per cent of the budget is identification, then halving the cost of an identification event is worth more than any conceivable improvement to every other line combined. That means better front-end rejection, cheaper thermal sensing, and inference that terminates early on easy cases—none of which is glamorous and all of which is decisive.

And it gives the **legal argument a quantity**. The companion paper argues that an emplaced offline system is better described by the law of mines than by the law of weapons under human operation. This section supports that from the engineering side and adds something the legal analysis cannot supply on its own: the two are separated by a measurable expenditure, so the question of how much discrimination the law should require is, in this class of system, answerable as a number of watt-hours. A rule that demanded identification before engagement would be demanding 830 mWh per emplaced-day. That is a rule with a price, and the price is small—which is the strongest argument available that the rule is a reasonable one to impose.

22. What none of this solves

Three limits, and the Part should be read against them.

Harvesting and scavenging both scale with something outside the operator’s control—weather in the first case, the intensity and location of fighting in the second. Neither yields a guaranteed budget. A system that must be certain of its energy for thirty days needs the primary-cell baseline at §19.2 and, above that baseline, the exchange route at §20, which means it needs a corridor a logistics robot can survive. Closure reduces the number of parties who can defeat the mission. It does not reduce it to zero, and the remaining party is the adversary.

The arithmetic is a design calculation. Every figure at §18.2 and §19.1 is derived from published component classes and stated assumptions. None is a bench measurement. The most likely place for it to be wrong is the tier-one and tier-two discrimination performance, because the whole budget is sensitive to the event rate and the event rate depends on how well the cheap tiers reject things that do not matter.

The energy result does not make the system lawful, useful or wise. It establishes that a thirty-day armed emplacement is not prevented by its power budget, which is a narrower claim than it may appear and is deliberately narrow. Whether such a system should be built is the subject of Part VII and of the companion paper, and nothing in this Part bears on it.

Part VII—What Closure Costs

23. The risks, stated at full strength

This section aims to set out what is lost by the architecture specified in Part V, without mitigation and before any answer is offered. Our purpose is to establish the cost honestly, because a proposal of this kind is most likely to fail not on its arithmetic but on a consequence its author found convenient to treat briefly.

Seven costs, in descending order of seriousness.

One—a system that cannot be reached cannot be stopped. At posture 0 there is no channel by which an operator can intervene between emplacement and expiry. If the tactical situation changes, if friendly forces enter the sector, if the emplacement was made on a mistaken picture, or if the system is behaving in a way its designers did not anticipate, nothing can be done until someone physically reaches it. This is the gravest cost in the list. It is not reducible to a probability, because the cases in which it matters are precisely the cases nobody predicted.

Two—a system with no update path cannot be corrected. A defect discovered on day three persists to day thirty. Part V §16.2 concedes that the platform’s discrimination will degrade in relative terms across the emplacement, because the adversary adapts and a fixed model does not. Everything the operator learns during those thirty days is unavailable to the system that most needs it.

Three—closure removes the audit trail. A supervised system generates a record: what it saw, what it decided, when, and on what basis. A closed system at posture 0 or 1 generates that record only on board, where nobody can read it until the platform is recovered, and not at all if the platform is destroyed. The consequence is that a failure may be invisible in exactly the way Part III said dependency failures are not: nobody learns that something went wrong, and so nobody corrects the class of error. This has legal consequences which are not this paper’s subject and which the companion paper takes up. It has an engineering consequence which is: the fleet does not improve, because the fleet is not told.

Four—capture. A platform left in contested ground for thirty days may be taken intact. Part V §16.5 treats keys, but the exposure is wider than keys: the adversary obtains the sensor suite, the inference model, the discrimination logic, the release criteria, and—most valuable of all—the knowledge of what the system does not detect. A captured platform is a specification of how to defeat every other platform of its type.

Five—a private link is only as good as its key management, and key management is where closed systems fail. Pre-shared material must be generated, distributed, loaded, tracked, revoked and replaced, by a formation under pressure, at scale, without the certificate infrastructure that makes this tractable in the commercial world. The failure mode is not cryptographic but administrative, and the classic one: the same key on too many platforms because that was easier.

Six—displacement, which is the failure that defeats the safety argument rather than the mission. Part V §16.2A sets this out: a platform moved by blast, subsidence, heave or a vehicle strike continues to enforce its surveyed sector over ground nobody surveyed. It is listed here because it is the only risk in this list that turns a mitigation into a hazard—every other entry degrades the system, and this one silently redirects it. §16.2A’s answer, displacement detection by invariant inter-feature geometry gating release, is a requirement of the specification and not an option.

Seven—a closed system forfeits genuine benefits, and they should be named rather than dismissed. Off-platform inference can be better than on-platform inference, because it is not power-constrained. A shared picture assembled across a theatre is better than one assembled across a group of six emplaced platforms. Vendor updates do fix real defects. The argument of this paper is that these are worth less than

the dependency they cost, not that they are worth nothing, and a commander who concludes otherwise for a particular mission has not made an error.

24. Which mitigations actually work

Each is assessed for what it closes and what it leaves open. Several are partial and are described as partial.

Physical geometry rather than software boundary—Part V §16.3. Closes: engagement outside the surveyed sector, robustly, against fault, corruption and intrusion alike. Leaves open: everything inside the sector, which is where the hard cases are. This is a strong measure with a narrow scope, and its strength is that it does not depend on the computing system being correct.

Physical safing not under software control—§16.4. Closes: unintended function of the weapon by any software cause. Leaves open: intended function on a wrong decision. It addresses the failure mode people fear and not the one that is more likely.

Per-platform keys with tamper erasure—§16.5. Closes: propagation of a single compromise to a fleet. Leaves open: the compromise of the captured platform itself, and all of risk four above except the keys. Erasure protects the network, not the secret of how the system works.

Tiered sensing—§16.1. Closes, incidentally, part of risk three: a tiered system's on-board record is small enough to be retained in full for the whole emplacement, because it records events rather than a video stream. Recovery of the platform therefore recovers a complete account. This does not help if the platform is destroyed.

Standardised exchange interfaces—§20. Closes: part of risk two, because a logistics robot delivering a pack can also carry a physical medium. An update route that requires physical presence is slow, and it is not nothing, and it does not fail the withdrawal test because the party carrying it is the operator.

Two mitigations deserve their own treatment because they do more work than the rest, and because the field has under-used both.

25. Receive-only, which is the most valuable design choice available

Posture 1 recovers most of what posture 0 surrenders at very close to zero cost, and the reason is asymmetric: **direction-finding locates transmitters**. A receiver that never transmits presents the same electromagnetic profile as a system with no link at all, while remaining commandable.

What that buys, against the risks at §23:

It substantially answers **risk one**. A system that can be told to stop can be stopped. The commander who discovers that friendly forces are entering the sector, or that the emplacement was made on a mistaken picture, has a channel. The recall problem is not eliminated—see the limits below—but it moves from impossible to unreliable, and that is a large move.

It partially answers **risk two**. Parameters can be changed within limits: sector boundaries narrowed, readiness levels altered, authority extended or curtailed, the system placed in a safe state pending physical attention. What cannot be delivered this way is a new model, because the bandwidth and the verification requirement are of a different order.

It does not answer **risks three to six** at all, and it should not be credited with doing so.

Three limits, and they are real.

No confirmation. The operator cannot know the command was received. The design must therefore be structured so that the *absence* of a message is the safe state rather than the dangerous one—which is precisely what expiry into inertness at §26 achieves, and it is why the two mitigations belong together.

Jamming denies the channel. A receiver can be denied even though it cannot be located, and an adversary who suspects a receive-only system exists can jam the band cheaply and without needing to find anything.

The answer is partial: multiple bands, a low-rate waveform with high processing gain, long integration and a message that can be repeated for hours because it is a few bytes.

A receiver is a surface. Anything that listens can be sent something, and the authentication of commands becomes the single point on which the whole benefit rests. This returns to risk five, and the reason §16.5 insists on the key ceremony.

25.1 Tiering the receiver, which the paper should have done from the start

The budget at Part VI §18.2 assumes a single duty-cycled receiver listening one second in sixty at 30 mW, giving 12 mWh per day. That figure applies the tiering principle to sensing and then abandons it at the antenna, and correcting the inconsistency improves the result by an order of magnitude.

Wake-up radio is a mature technique with its own survey literature.⁶³ Dedicated wake-up receivers operate in the microwatt class: reported designs draw a few microamperes, and average consumption of about 3 μ W has been demonstrated at a wake-up packet detection latency of tens of milliseconds.⁶⁴ The architecture is two-stage—an always-listening wake-up receiver detecting an addressed preamble, gating a capable main receiver that is otherwise unpowered—and it is precisely the tier-one-gates-tier-three pattern this paper argues for in sensing.

Applied to the control channel: *[Derived, on the sources cited.]*

wake-up receiver, 3 μ W continuous = **0.07 mWh per day** main receiver, 100 mW for 10 s on each of three commands a day = **0.83 mWh per day total approximately 0.9 mWh per day**, against 12 mWh for the flat duty-cycled design

The total budget moves only from 0.911 to 0.900 Wh per emplaced-day, because tier three dominates everything—which is itself a demonstration of the paper’s central point. What changes materially is the price of recall, which falls from 1.3 per cent of the budget to **0.1 per cent**. Recall is not merely cheap. It is very close to free, and the discussion of autonomous ground systems is conducted as though it were expensive.

Two consequences, and the second is the more interesting.

Sensitivity is the real constraint, not power. Reported wake-up receivers achieve sensitivity of the order of -90 dBm,⁶⁵ which is well short of what a sensitive tactical receiver achieves and therefore imposes a link budget the design must close. At five to fifteen kilometres this is a question of transmit power and antenna gain at the far end rather than of anything on the platform.

Which means the emission cost of recall is borne by the control station, not by the emplaced platform. The station transmits hard; the platform stays silent and simply listens well enough to hear. That asymmetry is favourable in a way that should be stated plainly: the station is to the rear, is already emitting for other reasons, can be mobile, and can transmit through a remote or decoy antenna at a location that is not its own. The exposure created by making the system recallable can therefore be placed somewhere it is cheap, rather than on the thing that is trying not to be found. No part of the recall channel requires the emplaced platform to reveal itself.

Set against the limits above, the cost of posture 1 over posture 0 is a receive chain drawing about 0.9 mWh per day—one part in a thousand of the budget—with the associated emission borne elsewhere. **It is by an enormous margin the cheapest risk reduction in the specification**, and that it does not feature in a debate generally conducted as though the choice lay between full supervision and none is the most useful observation in this Part.

⁶³Rajeev Piyare and others, ‘Ultra Low Power Wake-Up Radios: A Hardware and Networking Survey’ (2017) 19 *IEEE Communications Surveys & Tutorials* 2117.

⁶⁴ibid; and see ‘Leveraging Energy Harvesting and Wake-Up Receivers for Long-Term Wireless Sensor Networks’ (2018) 18 *Sensors* 1578, reporting average wake-up receiver consumption of approximately 3 μ W at a detection latency of tens of milliseconds.

⁶⁵‘A Tuned-RF Duty-Cycled Wake-Up Receiver with -90 dBm Sensitivity’ (2018) 18 *Sensors* 86.

26. Expiry into inertness

Authority should be time-limited, and the expiry state should be inert rather than continued.

The argument is about which failure the design prefers, and it is a design decision rather than a discovery. It is also one the law has already made for the analogous device, which is worth knowing before making it again. Amended Protocol II requires remotely delivered mines to self-destruct within thirty days at 90 per cent reliability, with a back-up self-deactivation leaving no more than one in a thousand functional after 120 days.⁶⁶ That is a two-layer design—a primary mechanism with a stated reliability and an independent backstop with a stated residual—and it is a better specification than anything this paper had arrived at unaided. It should simply be adopted: a primary expiry at the emplacement horizon, and an independent second mechanism, on a different physical principle, bounding the residue. A system whose default on losing contact, on exhausting its energy reserve, or on reaching its authority horizon is to *continue* has been designed so that neglect is indistinguishable from intent. A system whose default is to *stop* has been designed so that the dangerous state requires a positive act.

Three properties of the implementation matter and each follows from a failure mode above.

Expiry must be enforced by the physical safing rather than by software, because a software timer shares a failure domain with everything else in the computing system.

The interval must be short enough to bound the consequence of losing the receive channel. If commands cannot be delivered, expiry is the only remaining limit, so the horizon should be set at the period the commander is content for the system to operate unreachable—which is a tactical judgement and should be an emplacement parameter rather than a design constant.

Extension must be a positive act. Where the receive channel is available, extension by command is the ordinary case. Where it is not, extension requires physical presence, which is expensive, and that expense is the point: it makes continued operation a decision somebody takes rather than a state that persists.

Taken together with §25, the pairing produces the property the architecture needs. The system can be stopped if a message gets through, and it stops anyway if none does. Both failure directions resolve towards inert. That is the strongest available answer to risk one, it costs 12 mWh a day and a mechanical timer, and it is available to any system on this architecture.

26A. The adversary’s counter-architecture

A design should be assessed against a competent opponent who knows how it works, not against one surprised by it. This section games the response, and it produces the paper’s most robust finding: the architecture resists counter-measure for a structural reason rather than a technical one, and the structural reason also identifies the one place it is genuinely weak.

26A.1 What does not work against it

Assume the adversary has read this paper and has captured a platform, so knows the sensing tiers, the postures and the authority envelope.

Jamming fails. There is no control link to sever. At posture 0 there is nothing; at posture 1 jamming denies recall, which harms the operator’s ability to *stop* the system and does nothing to prevent it functioning. This should be stated plainly because it is counter-intuitive: **jamming a receive-only holding system makes it more dangerous, not less.** An adversary who understands the architecture has an active reason not to jam it, which is an unusual property for a military system to have and is a direct consequence of designing the safe state as the default.

Cyber attack has almost no surface. No inbound network path, no update channel, no telemetry, no licence check, no vendor. The remaining surface is the command receiver, which accepts a few bytes on a pre-shared key, and physical access.

⁶⁶Amended Protocol II (n 61), Technical Annex.

Supply-chain and service coercion fail. There is no proprietor to pressure. The whole of Part III describes leverage that this architecture does not offer.

Attrition is expensive by construction. Each platform is cheap, static, concealed and individually of low value. Destroying one costs a munition and a detection, and detection is the hard part because the platform is not emitting.

26A.2 What does work

Three things, and they are the honest assessment.

Physical clearance. Find and destroy each platform by systematic search, as with any emplaced obstacle. This works and it is exactly what the system is for: it costs the adversary time, effort, munitions and exposure, which is the definition of terrain denial succeeding. A counter-measure that consumes the adversary's resources at a favourable exchange is not a defeat of the system.

Bypass. Route around the denied ground. Also the system working, provided the ground denied is ground that matters. This puts a constraint on employment rather than on architecture: emplaced denial is worth its cost only where the terrain does not permit cheap bypass, which is a siting judgement and should be treated as one.

Attacking the cycle rather than the platform. This is the real vulnerability and it is not on the platform at all. The architecture requires emplacement, occasional energy exchange, and recovery. Each is a moment when a human or a logistics robot is present, moving, and locatable—and the exchange interval is a *pattern*, which is precisely what an adversary is good at detecting. **The system is at its most survivable in the state it spends 99 per cent of its time in, and at its most vulnerable in the transitions.**

26A.3 The design consequence

That last point inverts the usual optimisation and produces a specific recommendation.

Effort spent hardening the emplaced platform buys little, because the emplaced platform is already hard. Effort spent **lengthening the interval between transitions** buys a great deal, because it reduces the number of exposures directly. Every result in Part VI now reads differently: the energy work is not primarily about making the system cheap to run, it is about making it possible to leave the system alone for longer, and the operational value of the primary-cell baseline at §19.2 is not that it saves energy but that it removes a scheduled visit.

This also sets the correct target for the exchange model at §20. The question is not how efficiently energy can be delivered but how rarely it needs to be, and a design that halves the delivery frequency is worth more than one that halves the delivery cost.

27. What remains unmitigated

Stated plainly, because the mitigations above should not be read as disposing of the list.

Risk four is not mitigated in any serious sense. A captured platform discloses how the system works and how it can be defeated, and no measure in this paper prevents that. The only real answers are to design for it—assume the adversary has the specification, and do not rely on secrecy of method for effectiveness—and to keep the value of any single platform low enough that its loss is tolerable. Both are sound and neither is a mitigation.

Risk six is a trade and not a problem to be solved. Closure costs capability. This paper argues the trade is favourable for the holding mission and unfavourable for others, and Part VIII draws that line.

Risk three is the one this paper is least able to answer and the one that matters most beyond engineering. A system that acts without generating a contemporaneous record accessible to anyone answerable for it raises a question that engineering cannot settle, and the companion legal paper begins precisely where this section ends.

27A. What this makes the system, doctrinally

The paper has argued the architecture in engineering terms. It should be stated what the result is in the vocabulary of the discipline that would employ it, because the answer is not “a robot” and the difference is doctrinally consequential.

Countermobility doctrine treats emplaced effects as **obstacles**, and plans them by *obstacle intent*, which doctrine decomposes into three components: the **target**, the **obstacle effect**—disrupt, turn, fix or block—and the **relative location**.⁶⁷ Denial measures are the wider category: action to deny an enemy the use of space, personnel or facilities, by destruction, removal, contamination or obstacle construction.⁶⁸ The taxonomy is mature and the planning discipline around it is among the most worked-out parts of engineer doctrine.

Every obstacle in that taxonomy shares a property so universal that doctrine has never had cause to state it. **All three components of obstacle intent are fixed when the obstacle is emplaced.** A minefield laid to turn cannot be made to block; wire laid to disrupt cannot be re-tasked to fix; a demolition prepared to deny a crossing does one thing to one target in one place. Changing any component means physically re-working the obstacle under whatever conditions then obtain, which is why obstacle plans must anticipate the scheme of manoeuvre and why getting one wrong is expensive to correct. Lanes and gaps exist because opening and closing are hard, and the engineer effort and risk in each is a standing planning constraint.

A closed holding system at posture 1 is the first obstacle for which **two of the three components are mutable after emplacement.** The precision matters and it strengthens the claim rather than weakening it. *Relative location* is fixed: the platform sits where it was put, and moving it is exactly as hard as moving any other obstacle. *Target* and *obstacle effect* are not. The sector can be narrowed or widened within the arcs, the release criteria tightened or relaxed, readiness raised or lowered, the system placed inert and revived—by a command carried on a channel costing, per §25.1, about one part in a thousand of its energy budget, and radiated by a station to the rear rather than by the obstacle itself. Nobody has to go forward. Nothing has to be dug.

That an obstacle can change its effect while remaining where it is has no precedent in the taxonomy, and it is not a marginal capability. Effect is the component that determines how the obstacle relates to the scheme of manoeuvre, and the component most likely to be wrong when the plan is made, because it depends on what the enemy does.

Three consequences, and they are doctrinal rather than technical.

An obstacle that can be opened and closed is a different planning object. A gap through a minefield must be planned, marked, guarded and closed by engineers under fire. A gap through a robotic obstacle is a parameter change. The counter-attack that must pass through its own obstacle belt—a recurring and expensive problem—becomes a scheduling question rather than an engineering task.

The obstacle plan need no longer anticipate the whole scheme of manoeuvre. Its central difficulty is that it commits early and cannot be revised. A mutable obstacle can be committed provisionally and adjusted as the situation develops, which reduces the cost of a wrong initial judgement to approximately nothing.

And it makes the recall channel a doctrinal requirement rather than a safety feature. §25 argues for the receive channel on grounds of risk. This is the positive case for the same design: the channel is what makes the object an adjustable obstacle rather than a fixed one, and without it the system is a mine with better sensors. That is a stronger argument for building the channel than any of the safety arguments,

⁶⁷US Army, *FM 90-7: Combined Arms Obstacle Integration*, on obstacle intent as target, obstacle effect and relative location, and on the four obstacle effects; and US Army and US Marine Corps, *ATP 3-90.8 / MCWP 3-17.5: Combined Arms Countermobility Operations* (September 2014), on obstacle intent and obstacle employment principles.

⁶⁸On denial measures see US Army, *FM 5-102: Countermobility*, ch 7 (‘Denial Operations’) <https://www.globalsecurity.org/military/library/policy/army/fm/5-102/Ch7.htm> accessed 5 August 2026. On the contemporary framing of terrain denial and the observation that Army doctrine addresses denial of terrain to the enemy without using the term as a category of fires, see Lieber Institute, ‘Terrain Denial Fires on the Modern Battlefield’ <https://lieber.westpoint.edu/terrain-denial-fires-modern-battlefield/> accessed 5 August 2026.

because it appeals to what a commander wants rather than to what a lawyer requires—and it happens to deliver the same thing.

The claim is bounded in one respect that should be stated. Mutability is only as good as the delivery of the command, and §25 sets out why delivery is unreliable under jamming. An obstacle whose intent can usually be changed is not the same as one whose intent can always be changed, and an obstacle plan should be sound on the assumption that no command gets through. What mutability buys is an improvement on that plan, not a substitute for it.

Part VIII—The Boundary, and Conclusions

28. Where closure is the wrong answer

This paper has argued for a design. A design that is recommended everywhere has not been specified, so the boundary matters as much as the case.

Closure is wrong where the mission depends on information that did not exist at emplacement and cannot be obtained on a private link. Targeting against a picture assembled at theatre level, coordination with formations that do not share keys, and any task requiring current external data are cases where the dependency buys something the operator cannot otherwise have. The correct response is to accept the dependency and to know its price, which the audit method at Part III §9 supplies.

Closure is wrong where the cost of a wrong engagement exceeds the cost of the emission that would prevent it. This is a proportionality judgement rather than an engineering one, and it should be made explicitly rather than being settled by an architectural default. Where the consequence of error is high—populated ground, contested identification, proximity to friendly or neutral parties—the emission required to keep a human in the loop is cheap at the price.

Closure is a poor fit for logistics and assault, though less poor than an earlier formulation of this argument assumed, and the correction is worth recording. The objection to closure for those missions is that they need live coordination. They do, and a private mesh at posture 2 or 3 supplies it. What they genuinely need and closure genuinely burdens is *routing against conditions that change faster than a cached picture*, and *coordination across formations that were not keyed together in advance*. Those are real constraints. They are narrower than “these missions need networks”, and stating them narrowly is the correct outcome of separating the two axes.

Closure is wrong where the target is fast, and there is a combat-tested system that shows exactly where the line falls. Ukraine’s Sky Sentinel is a static, armed, autonomously operating turret: a heavy machine gun on a 360-degree mount, cued by radar, which detects, locks, tracks, computes and fires without a soldier operating it, tested in combat, at a unit cost of about USD 150,000, and credited with engaging targets at speeds up to 500 mph and objects five times smaller than a Shahed.⁶⁹ It is the closest fielded thing to the architecture specified in Part V, and it is built on the opposite sensing philosophy—active radar, continuous readiness, and the power budget both imply.

The difference is not a disagreement about design. It is the target. **The speed of the thing being engaged sets the sensing modality, and the sensing modality sets the emission and the power budget.** An air target closing at 120 to 500 mph must be detected at kilometres and engaged within seconds, which requires active sensing and forecloses everything Part VI depends on. A dismounted man at walking pace can be detected at thirty metres by a passive geophone and engaged over tens of seconds. The architecture of this paper is available for ground denial *because ground targets are slow*, and Sky Sentinel is the proof of what the alternative costs when they are not.

⁶⁹David Hambling, ‘Ukraine Turns To AI-Controlled Guns To Stop Russian Shahed Drones’ (*Forbes*, 27 May 2025) <https://www.forbes.com/sites/davidhambling/2025/05/27/ukraine-turns-to-ai-controlled-guns-to-stop-russian-shahed-drones/> accessed 5 August 2026; *Militarnyi*, ‘Sky Sentinel: Ukraine to produce anti-aircraft turrets for intercepting drones and cruise missiles’ <https://militarnyi.com/en/news/sky-sentinel-ukraine-to-produce-anti-aircraft-turrets-for-intercepting-drones-and-cruise-missiles/> accessed 5 August 2026. Performance and cost figures are as reported by the developer and by a party to the conflict, and are treated on the footing set out in the note on sources.

That yields a clean statement of the boundary. Closure and a microwatt floor are available where the engagement timeline is measured in tens of seconds and the target is surface-bound. They are unavailable against air, against fast vehicles at range, and against anything requiring detection beyond the reach of passive sensing. A commander who needs both is buying two systems, not one.

And there is a fielded static armed sentry which shows what the conventional architecture buys in this exact role. The SGR-A1, developed by Samsung Techwin and Korea University and reported in the Korean demilitarised zone from 2010, mounts a 5.56 mm light machine gun with an optional grenade launcher, uses infrared and thermal imaging with pattern recognition to detect and track at ranges reported up to three kilometres, and issues voice warnings before engaging.⁷⁰ Unit cost was announced at about USD 200,000 in 2006; the number deployed is not public.

It is the closest fielded system to the design case in Part V, and it is built the other way round on every axis this paper separates. It is a fixed installation with mains power, so the emplaced-day budget does not constrain it. It uses thermal imaging continuously rather than as a woken tier, which is what buys the three-kilometre detection range against this architecture's tens of metres of passive cueing. And it sits on a permanently monitored, marked and lawfully closed boundary, which is the environment in which the durable propositions of §22 hold indefinitely.

The comparison is the boundary stated in hardware rather than in argument. Where power is available, the ground is permanently controlled, and detection at kilometres is the requirement, the conventional architecture is correct and this paper's is not. Where the platform must be carried forward, concealed, powered from what it carries, and left in ground nobody controls, the trade reverses. Neither system is a better version of the other; they answer different questions, and the questions are separated by whether there is a wall to bolt the thing to.

Closure is wrong where it is used to justify what should not be built. A system architected so that nobody can reach it is not made acceptable by being efficient. Part VII §27 does not resolve into a mitigation, and the paper does not pretend otherwise.

29. Conclusions on each element

The two-axis separation holds and the paper's most portable result. Closure and emission are independent properties and they have been treated as one. The withdrawal test—does the mission fail if a party other than the operator withdraws a service—can be applied to any specification by enumeration, and it produces a count. The finding that follows is that the choice the field believes it faces between capability and survivability is not one, because the capability at issue is coordination and coordination is available on a private mesh at no cost to closure.

The dependency audit establishes shape rather than magnitude, and its most useful figure is a negative one. Of eleven classes of service consumed by a representative stack, ten are discharged during the mission and between four and seven are supplied by another party. Eight of the ten can be lost without an adversary doing anything. One class is worse than a dependency and was not in the first version of the audit: an unacknowledged vendor relationship, of the kind the reporting around the platforms demonstrated at *Golden Dragon 2024* describes, removes the party who could withdraw and puts nobody in their place—no contract, no support, no route to demand disclosure. Grey acquisition solves the withdrawal problem by making the sustainment problem total, which is an argument for closure that owes nothing to Part II. The exposure is therefore not principally a jamming problem, and hardening does not address a supplier's terms of service. The audited stack is a composite and no fielded product has been audited against its own specification, which is a limit of access rather than of method.

The emission budget is stated as scaling relationships rather than as ranges, and that is the strongest form available. Detection radius exceeds communication radius structurally. Duty cycle,

⁷⁰On the SGR-A1 see 'SGR-A1' <https://en.wikipedia.org/wiki/SGR-A1> accessed 6 August 2026, and IEEE Spectrum, 'A Robotic Sentry For Korea's Demilitarized Zone' <https://spectrum.ieee.org/a-robotic-sentry-for-koreas-demilitarized-zone> accessed 6 August 2026. Development is attributed to Samsung Techwin, now Hanwha Aerospace, with Korea University. Deployment numbers are not public and the programme is reported as classified, so the figures here establish the system's character rather than the scale of its use.

waveform and directivity act on different parts of an interceptor's problem, and a static platform can push all three favourably at once in a way a mobile one cannot. Persistence rather than power is the quantity that governs geolocation, which means designers optimising transmit power while accepting a continuous link have optimised the wrong term.

The metric proposal is a definition and stands or falls on usefulness. Endurance in drive hours measures locomotion, which is the whole budget of a moving platform and approximately none of a static one. Watt-hours per emplaced-day at a stated readiness level is proposed instead, with energy per square kilometre per day of denial above it. The second permits robotic and infantry terrain denial to be compared on one axis, and this paper does not make that comparison because it lacks the sustainment figures to make it honestly.

The architectural inversion is the paper's central engineering claim, and it now rests on a fielded comparison rather than only on a calculation. REMBASS operates unattended for thirty days at a thousand detected events per day, against approximately eight hours for the ground surveillance radar the programme literature compares it with: the same task, the same era, a factor of about ninety, and architecture rather than chemistry as the difference. The calculation at §18.2 reaches the same conclusion by a different route. **This has not been built and the arithmetic is not a measurement.** Its most likely failure point is the discrimination performance of the cheap tiers, because the budget is dominated by tier three and tier three is driven by the event rate rather than by tier three's own power draw: across the published range of 3 to 20 W for an uncooled imager and an edge inference module, the thirty-day requirement moves only from 10 Wh to 52 Wh, and the harvesting surface from 0.9 Wp to 4.7 Wp. The design is insensitive to a factor of seven in the largest line of its own budget, which is a consequence of the duty cycle and is the most reassuring property it has.

The interlock is the result the author would most want tested. Ambient harvesting is dismissed for this application on yield and signature, and the dismissal is correct against a vehicle-class load and wrong against an inverted one—the December requirement at Ukrainian latitudes falls to a surface smaller than a sheet of A5. The cold-charging objection, which is the serious one, is answered not by heating but by not charging: a lithium thionyl chloride primary baseline carries the 76 mWh per day that must never fail, and a small harvested rechargeable buffer carries the bursty remainder, so that cold degrades readiness rather than causing failure.

Energy is better understood as a currency than as a constraint, and what it buys is discrimination. Re-read as a price list, the budget puts identification and release judgement at 91 per cent of the total, recallability at 1.3 per cent, and the weapon itself at 0.5 per cent. A mine is the zero-energy limit of the same design—it spends nothing and exercises no judgement—so the two are one class of object separated by a measurable expenditure. Amended Protocol II's obligations can then be priced: self-deactivation and detectability cost approximately nothing, and discrimination, which the law never required of mines because no unsensing device could supply it, is the one line that costs anything. That is the most portable result in Part VI, because it converts “how much discrimination should the law require” into a question with a number attached, and the number is small.

Stasis has design consequences the field has not worked out, and one of them is a defect in this paper's own safety argument. A static sensor subtracts its background from one survey, needing no simultaneous localisation and mapping and no continuous ranging, which takes lidar out of both the emission and the power budget and is part of why §18.2 closes. A static platform's cost is what it pays to do nothing, which inverts the morphology ordering. And a displaced platform enforces its surveyed sector over unsurveyed ground—the geometric constraint does not fail, it relocates, and every argument here that treats physical geometry as reliable is exposed to it. Self-localisation by inter-feature geometry answers it: pairwise ranges and angles among fixed features are invariant under displacement of the observer while absolute bearings are not, so the platform can distinguish “I moved” from “the world changed” using nothing external, which preserves closure where an inertial unit or a positioning fix would not. Detection of displacement must gate release rather than re-deriving a sector, and the same self-survey at power-on permits robotic emplacement, which attacks the transition §26A identifies as the architecture's real vulnerability.

The system is best understood doctrinally as a mutable obstacle, and that is the paper's

contribution to the discipline that would employ it. Countermobility doctrine plans obstacles by intent, and every obstacle in the taxonomy has its intent fixed at emplacement—a minefield laid to turn cannot be made to block without physically re-working it under whatever conditions then obtain. A closed system at posture 1 can be re-tasked by a command costing one part in a thousand of its energy budget and radiated from the rear. Gaps and lanes become parameter changes rather than engineer tasks under fire; the obstacle plan need no longer commit to the whole scheme of manoeuvre in advance. This also supplies the positive case for the recall channel: without it the object is a mine with better sensors, and with it the object is something the taxonomy has never contained. Mutability is bounded by delivery, and an obstacle plan should remain sound on the assumption that no command arrives.

The energy cost of holding ground is set by the false-alarm rate, not by enemy activity. The escalation rate is the output of a detection cascade, $N = T \cdot P \cdot P + F \cdot q$, and wherever tier-one false trips greatly exceed genuine events—the normal condition for an unattended sensor, since the environment is busy and the enemy is not—the second term dominates. A sector a company crosses daily costs almost what an empty one costs; a sector in high wind or near a road costs several times more. Two design results follow: tier-two false-alarm rejection is the highest-leverage parameter in the system, and detection range is coupled to energy cost through the false-alarm rate, so the optimum tier-one sensitivity sits below what the hardware permits. Neither parent literature contains that result, because the unattended-sensor field optimises detection subject to power and the ground-robotics field does not treat detection as an energy question.

The reporting trade is decidable and nobody has stated the rule. A holding system’s intelligence product and its survivability compete for the same account. Report when the expected value of the information exceeds the expected cost of the disclosure; set the thresholds at emplacement, alongside the authority envelope, because the commander who emplaces the system knows what would change their mind; and batch what is retrospective, because most of the value of these observations is pattern-level rather than immediate.

The architecture resists counter-measure structurally, and its weakness is not on the platform. Jamming a receive-only holding system denies recall and does not prevent function, so it makes the system more dangerous rather than less—an adversary who understands the design has an active reason not to jam it. What works is physical clearance and bypass, both of which are the system succeeding, and attacking the emplacement, exchange and recovery cycle, which is where humans and logistics robots are present, moving and locatable on a detectable interval. The design consequence inverts the usual optimisation: hardening the emplaced platform buys little because it is already hard, and lengthening the interval between transitions buys a great deal. The energy work is therefore not principally about running cheaply. It is about being able to be left alone for longer.

Exchange rather than recharge is the delivery model, and it requires nothing new. The fleet conducting most forward logistics in some formations already carries cargo of the required mass. What is needed is a standardised pack format, an interface a manipulator can operate blind in mud, and a platform that runs on its primary cells throughout the exchange.

Scavenging is speculative and is the paper’s least supported proposal. The source is real and unusually well suited—co-located with the mission and counter-cyclical to every other logistics curve—and the arithmetic survives a recoverable fraction an order of magnitude below the assumed one. Four engineering problems stand in the way and state-of-health assessment on unknown damaged cells is the hard one. §21.3 states the experiment.

30. What would change the analysis

A coordination function of real operational value shown to be unobtainable without a third-party service. This would defeat the central separation and the paper’s argument would have to be reduced to a claim about cost rather than about architecture. The author has not identified one; that is not the same as there being none.

A measured event rate materially above the assumed one. The budget is sensitive to it and nothing else. A field measurement of how often a tiered front end in real terrain escalates to the imaging tier is the

single most valuable experiment anyone could run against this paper, and it is cheap.

A measured cost of an identification event materially above 830 mWh. The exchange rate at §21A rests on that figure, and the whole reframing of energy as the price of discrimination moves with it. A thermal imager whose settling time is longer than assumed would change the number without changing the argument's shape; a factor of ten would change both.

Demonstration that receive-only channels cannot be made reliable enough to be relied on for recall. Part VII §25 and §26 carry most of the paper's answer to its gravest risk. If the channel cannot be delivered under jamming often enough to matter, then expiry into inertness is doing all the work alone, and the horizon would have to be set far shorter than a month—which would substantially reduce the value of the whole architecture.

31. Summative remarks

A machine gun on a tracked chassis held a position for a month, and the discipline that built it has no metric that describes what it did. That is the observation this paper started from and it turned out to be a more useful one than it first appeared.

The reason there is no metric is that the field measures drive hours, and drive hours measure locomotion, and the defining property of holding ground is that nothing moves. Once the right quantity is asked for, the architecture that supplies it turns out to exist already in an adjacent discipline that this one does not cite. And once the platform is architected that way, three problems that were being treated separately—dependency, detectability and energy—turn out to be one problem, because the design decision that closes the system is the decision that quiets it, and the decision that quiets it is the decision that makes its energy budget small enough to be met by a surface the size of a sheet of paper.

None of that has been built, and this paper is careful throughout about the difference between an arithmetic that closes and a result that has been measured. What is offered is a specification precise enough to be argued with, a budget with every input stated, and, in each case where the argument could be wrong, the experiment that would show it.

What the paper does not offer is any view on whether such systems should exist. The energy result establishes that a thirty-day armed emplacement is not prevented by its power budget. That is a narrower finding than it may appear, and it is deliberately narrow, because the question it leaves open is not an engineering question. A companion paper takes it up, and it begins at Part VII §27, with the observation that a system which acts and keeps no record anyone can reach has removed the thing every account of responsibility depends on.

The derivations in Part VI are reproducible from the figures stated in the text. The Appendix sets out the build and the measurements that would move each claim from an arithmetic to a result.

Appendix—What Would Have to Be Built

This paper makes a design argument and reports an arithmetic. Neither is a measurement. This Appendix states what would have to be constructed and measured to move each claim from the first category to the second, in the order of how much the paper depends on it, and it is written so that another group could run these tests without reference to the author.

The convention used throughout the paper is repeated here because it governs how the Appendix should be read. *Measured* means a published figure from the technical literature or a manufacturer's data. *Reported* means a statement by a party to the conflict about its own performance. *Derived* means the author's arithmetic from the first two, with the inputs stated at the point of use.

A.1 The experiment that matters most

The event-rate measurement. The budget at Part VI §18.2 is dominated by tier three, and tier three is driven by how often the cheap tiers escalate. Everything else in the energy argument is comparatively insensitive. The single most valuable experiment against this paper is therefore also the cheapest one.

Build: a tier-one and tier-two front end only—geophone, small acoustic array, magnetometer, passive infrared, and the escalation logic—with no imaging tier and no weapon. Log every escalation decision with a timestamp and the triggering modality.

Emplace: in representative terrain, for a representative period, at several sites differing in traffic, weather and proximity to a road.

Measure: not “escalations per day”, which does not transfer between sites, but the three terms of the cascade at §18.3 **separately**—the rate of genuine events entering the sector, the tier-one false-trip rate, and the fraction of tier-one false trips that tier two passes. Record each against ground truth, and repeat across threshold settings so that the sensitivity-versus-false-alarm curve can be drawn. Those three transfer to a different site, season and threshold; the composite does not.

Decides: whether the budget at §18.2 is realistic. Ten escalations a day gives 0.91 Wh per emplaced-day. A hundred gives 8.4 Wh, which the interlock at §19 still supports. A thousand gives 83 Wh, and the architecture is in trouble at the harvesting sizes claimed. The paper’s arithmetic assumes ten and the assumption is not defended by evidence, only by plausibility.

This experiment requires no weapon, no autonomy, no classified material and no unusual funding. That it has apparently not been run and published is the largest single gap in the public record on this subject.

A.2 The power measurements

Tier-one floor, at a stated detection range. Measure the continuous draw of an integrated seismic and acoustic front end with wake-on-event, including the microcontroller and its real-time clock, at the sensitivity required for reliable detection at the design range. The paper assumes 500 μ W for the tier. The range is the part to hold fixed while measuring: published field trials give reliable footstep detection at about 30 m and vehicle detection tested to 30 m, with more favourable conditions reported to 94 m, and the figure is a property of the soil as much as of the sensor, so it must be measured on ground representative of the intended siting. *Decides:* whether the always-on floor is where the unattended-sensor literature implies it is when a detection requirement is imposed on it.

Tier-three cost per event. Measure the energy of a complete escalation: thermal imager settling, frame capture, inference, decision, return to sleep. The paper assumes 10 W for thirty seconds, which is 0.083 Wh. Settling time is the term most likely to be underestimated, because a cooled imager is not instantly available and an uncooled one may need longer integration.

Receiver cost, and command delivery under jamming, which matters more. Part VII §25.1 replaces the flat duty-cycled receiver with a two-stage wake-up architecture at approximately 0.9 mWh per day, on published figures for wake-up radios. Two things need measuring and the second matters far more than the first. Measure the average draw of the two-stage chain at a detection latency adequate for the mission. Then measure **command delivery probability under representative jamming**, at the link geometry of §16 and at the sensitivity a wake-up receiver actually achieves, since reported designs sit around -90 dBm. That second measurement carries the recall argument at §25, the doctrinal claim at §27A that the system is a mutable obstacle, and the safe-failure argument at §26. If commands do not arrive often enough, expiry into inertness is doing all the work alone and the emplacement horizon must be set far shorter than thirty days.

A.3 The harvesting and cold tests

Yield at size and season. Measure delivered energy from a concealed, terrain-matched harvesting surface of the order of 0.015 m², at 50°N, in December, over at least thirty consecutive days including snow. The

paper derives approximately 370 mWh per Wp per day from an annual figure and a stated seasonal fraction with a 50 per cent derating. That derivation is the weakest link in §19.1 and a direct measurement would replace three assumptions at once.

Signature. Measure the optical, thermal and radar return of the harvesting surface against the surrounding terrain, at the ranges and sensor types of concern. The paper asserts that a surface of this size is concealable and does not demonstrate it.

Cold behaviour of the split architecture. Measure the primary-cell baseline and the harvested buffer through a full winter cycle, including the case where the buffer remains uncharged for an extended period. *Decides:* whether the graceful-degradation claim at §19.2 holds, or whether an uncharged buffer produces a failure rather than a lengthened response interval.

A.4 The scavenging feasibility study

Stated at §21.3 and repeated here in the form of a protocol.

Sample: recover expended first-person-view airframes from a representative area of ground over a defined period, by systematic search rather than by opportunistic collection, recording the search area and effort so that a rate can be computed.

Measure, in sequence: the fraction with a physically intact pack; the fraction of those which are electrically serviceable; the distribution of chemistry, cell count and capacity; and the time and energy required to assess state of health to a confidence sufficient to accept or reject.

Decides: the recoverable fraction, which the paper assumes at 10 per cent. The proposal survives 1 per cent and does not survive 0.1 per cent.

Note on safety: this protocol handles damaged lithium-polymer cells and should be conducted with the venting, separation and jettison provisions described at §21.2 in place from the outset rather than added after an incident.

A.4A The displacement-detection test

Part V §16.2A makes displacement detection a requirement of the specification and asserts that it is nearly free. Both parts need testing and the second is the one that would sink it.

Build: the tier-three imaging chain only, with a routine that extracts inter-feature geometry—pairwise ranges and included angles among fixed scene features—from a captured frame and compares it against a stored reference.

Measure: the smallest displacement reliably detected, in translation and in rotation, across scene types from feature-rich to near-featureless; the false-displacement rate under changing light, season, snow cover and partial obscuration, since a false positive reverts the platform to safe and a system that cries displacement is useless; and the energy and time added to a tier-three escalation.

Decides: whether the answer at §16.2A is a design or an aspiration. The claim is that the computation rides on imagery already captured and costs approximately nothing. If it needs its own sensor, its own frames or its own power budget, the claim fails and displacement detection becomes a line in the budget at §18.2 rather than a free rider on it.

A.5 The dependency audit, repeated properly

The audit at Part III is a composite assembled from published descriptions. The version worth having is one conducted against real specifications.

Method: for a specific fielded platform, obtain the full service enumeration—including entitlement checks and telemetry, which are the two classes least likely to be volunteered—and apply the withdrawal test at §6.1 to each.

The obstacle, and the finding in advance: this cannot be done from public documentation, and a purchaser who does not require the disclosure as a procurement condition cannot do it either. The method is sound and the data is withheld, which is a policy problem rather than a technical one and belongs to the companion paper.

A.6 What the paper got wrong in draft, recorded because it bears on the argument

Two corrections were made during preparation and both are instructive enough to report.

The paper originally argued for “offline” systems and the word was doing two jobs. It was used to mean independence from third-party infrastructure and to mean absence of emission, and the conflation led the argument to concede logistics and assault to the networked position on the ground that those missions require live coordination. They do require it; they do not require it from a third party. Separating the two axes recovered the concession and produced the withdrawal test, which is now the paper’s most portable result. The error was a definitional one and it had substantive consequences, which is the ordinary way with definitional errors.

The first version of the energy budget was computed against a load that included continuous imaging, on the unexamined assumption that a watching system watches. The unattended-sensor literature makes it plain that this is a design choice rather than a requirement, and correcting it changed the result by a factor of several hundred. The correction was available in an adjacent field’s ordinary practice and was missed because the two literatures do not cite one another, which is the specific gap this paper exists to close.

A.7 Supplementary material

The derivations in Part VI are reproducible in full from the figures stated in the text; no code is required and none is supplied, which is a deliberate consequence of keeping every calculation simple enough to be checked by hand. Every input figure is either cited to a source or marked as an assumption at the point of use. A reader who disagrees with an assumption can substitute their own and recompute, and §18.2 and §21.1 are written so that the sensitivity of the result to each assumption is visible without re-deriving it.
