

The Emplaced Machine

Energy, Detection, Recall and Control in Closed Ground-Denial Systems

Christopher I. V. Farmer, LL.M.

Independent researcher | Geneva, Switzerland

VERSION 2.12.5 | 17 OCTOBER 2026

ENGINEERING PREPRINT | NOT PEER REVIEWED

Suggested citation. C. I. V. Farmer, “The Emplaced Machine”, preprint, v2.12.5 (2026), <https://doi.org/10.5281/zenodo.23073940>.

Earlier title. Versions through v2.10.0 circulated as Holding Ground with Machines That Do Not Move: v1.0.0 (6 August 2026) was subtitled Why Endurance Is the Wrong Measure for Terrain Denial, and What a Thirty-Day Armed Emplacement Actually Costs in Watt-Hours, and later versions Endurance Metrics for Terrain Denial and the Watt-Hour Cost of a Thirty-Day Armed Emplacement. The two published legal companions refer to it, without a title, as the companion paper on closed-system ground robotics.

Revision record. The complete itemised history of corrections, narrowed claims and source upgrades appears at Appendix A.7. Version 2.12.5 restates as targets and plans reported figures that had been given as accomplished, withdraws figures not found in the sources cited for them, narrows three statements about the reported employment of ground robots, and corrects the earlier-title note and seven cross-references; it changes no model input, equation or numerical result. Version 2.12.4 updated the treaty analysis at §51, and the passages that depend on it, to the report the Group of Governmental Experts adopted on 4 September 2026, CCW/GGE.1/2026/3. Version 2.12.3 corrected one displayed rounded value.

Keywords. unmanned ground vehicles; terrain denial; emission control; unattended ground sensors; energy autonomy; closed systems; electronic warfare; military robotics; lidar; active sensing; point clouds; sensor fusion; authority envelope; modality substitution; range–energy trade

Licence. Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). Commercial use needs a licence from the author: civfarmer@gmail.com. The reproducibility archive (v2.12.2) remains under CC BY 4.0, as its README states.

Contents

Abstract	4
Extended summary	4
Note on sources and method	9
Note on the two legal companion papers and corrected figures	10
Part I—Introduction	11
1. Statement of the problem	11
2. Limitations of the current recommendation	11
3. Scope and approach	12
4. Contributions and prior work	13
Part II—The Two Axes: Closure and Emission Posture	19
5. Two senses of “offline”	19
6. Axis one: closure, and the withdrawal test	20
7. Axis two: emission posture	21
Part III—The Dependency Audit	22
8. Services consumed by a representative control stack	22
9. Limits of the audit	26
Part IV—The Emission Budget	27
10. Scope and limits of Part IV	27
11. The governing asymmetry	28
12. Detection and location	29
13. Comparison of emission postures	30
14. Reporting thresholds and the transmission decision	31
Part V—Specification of the Architecture	32
15. The holding platform	32
Part VI—Energy	58
16. Endurance metrics and the holding mission	58
17. The architectural inversion	60
18. The interlock	79
19. Energy delivery	81
20. Scavenging	82
21. The discrimination exchange rate	83
22. Remaining problems	87
Part VII—The Costs of Closure	88
23. The risks	88
24. Assessment of mitigations	89
25. The receive-only channel	90
26. Expiry into inertness	97
27. The adversary’s counter-architecture	98
28. Unmitigated risks	100
29. The denial test, and the safe-by-default emplacement	100

30. Doctrinal classification of the system	102
Part VIII—The Public Record	104
31. Statement of method	104
32. The funded public record	105
33. Ballistic computation as a catalogue item and a funded line	108
34. What is on sale to the public	109
35. The sensing principles, from the open literature	110
36. Atmospheric propagation: conditions and limits	113
37. Two programmes in precision small-arms fire, and where their public account ends	114
38. What the published record says about its own scope	115
39. The boundary of the record moves, and which way a rule-maker should err	116
40. What the compilation looked for and could not find	119
41. What the public record establishes, stated plainly	120
Part IX—The Function Chain	122
42. What this Part does, and the unit it uses	122
43. Waiting: unattended persistence with a live effect	123
44. Classifying: matching an observed signature against a stored profile	124
45. Selecting: engaging without a person designating the object	125
46. Accepting: autonomous engagement already fielded, already accommodated in policy	126
47. Re-tasking: changing an emplaced weapon's state from the rear	127
48. Supervising: what a person monitoring an unattended sensor field achieves	128
49. The finding: autonomy is a mode, and the safeguard is what was added	130
50. What this record does and does not contain	132
51. What a rule drafted now would reach	135
Part X—The Boundary, and Conclusions	138
52. Limits of closure	138
53. The state transition within the emplacement	140
54. Conclusions on each element	141
55. Sensitivity of the conclusions	144
56. Summative remarks	145
Appendix—Validation Requirements	145
A.1 The principal experiment	146
A.2 The power measurements	146
A.3 The harvesting and cold tests	148
A.4 The scavenging feasibility study	148
A.5 The survey-integrity test	148
A.6 The dependency audit, repeated properly	149
A.7 What the paper got wrong in draft, recorded because it bears on the argument	149
A.8 Reproducing the budget	157
A.9 The recall link and jam budget	157
A.10 Supplementary material and the reproducibility package	161
A.11 Active-sensor envelope and modality-substitution test	163
Reference List	166

Technical and scientific literature	166
Legal and policy materials	167
Government and programme records	169
Manufacturer, industry and press sources	171
Datasets, code and supplementary materials	174

Abstract

Endurance for a fielded ground robot is published as a drive-hour figure. The mission examined here—holding ground—is defined by not moving, so drive hours do not measure it. This paper proposes **watt-hours per emplaced-day at a stated readiness level** as the quantity in which an emplaced armed system should be specified, and computes it for a thirty-day emplacement from an eleven-parameter vector.

Three results follow. **First**, “offline” conflates two independent properties—independence from third-party services, and not emitting—and only the first is architectural; a withdrawal test separates them, and a dependency audit of a representative fielded control stack counts the parties whose withdrawal breaks the mission. **Second**, an emplaced armed platform should be architected as an unattended ground sensor that has effectors rather than as a vehicle that has stopped. The resulting budget is **1.22 watt-hours per emplaced-day**, against roughly 360 for a conventionally architected platform. **Third**, and the result that transfers outside engineering: **91 per cent of that budget is bought by identification and release judgement**. Holding ground without discriminating costs 110 milliwatt-hours a day; looking first costs about eleven times as much, and across the sampled uncertainty the multiple never falls below about six and a half. Discrimination is the expensive part, so its absence is what a market supplies cheaply.

The usable envelope of such a system is the intersection of five regions—where it can cue, where it can identify, where its surveyed geometry is still valid, where it is still authorised, and where the effector reaches—so a sensing range and an authority range are different quantities, and only the first moves when the sensor changes. **Modality substitution can expand sensing capability but may not expand human-conferred authority**, which requires the control to sit downstream of every perception path rather than inside one of them.

Two further design results follow. In the idealised processing-gain model, for a fixed occupied bandwidth and authenticated payload, each doubling of message duration adds approximately **3.01 dB** of processing gain; acquisition, coding and implementation losses make the tabulated delivery times lower bounds and the realised gain smaller. A **denial test** then shows that a stop-command channel is fail-dangerous under communications denial, whereas a periodically renewed permission-to-continue channel safes the platform after a bounded interval. The latter is **not “self-deactivation” as Amended Protocol II defines that term**: it does not irreversibly exhaust an essential component. It is an automatically functioning loss-of-authorisation safing mechanism.

Keywords. unmanned ground vehicles; terrain denial; emission control; unattended ground sensors; energy autonomy; closed systems; electronic warfare; military robotics; lidar; active sensing; point clouds; sensor fusion; authority envelope; modality substitution; range–energy trade.

Extended summary

The endurance of a fielded ground robot is a drive-hour figure: four to eight hours for small platforms, eight to sixteen for mid-sized logistics types, a stated next-generation ambition measured in days. The mission this paper examines—holding ground—is defined by not moving, and drive hours are the wrong measure for it. 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, set a target of contracting 25,000 vehicles in the first half of the year, and states the objective that all frontline logistics be performed by robotic systems. In the 3rd Assault Brigade, and in the contested cities of Pokrovsk and

Myrnograd, these platforms were reported in March 2026 to perform 80 to 90 per cent of logistics. The design case is a thirty-day emplacement—a stipulation taken from a treaty instrument, not a figure derived from an operational report.

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 energy figures throughout rest on eleven cascade parameters together with separately stated receiver, housekeeping and effect assumptions; all are component classes or stipulations rather than measurements, and Appendix A.2 sets out the single experiment that would close the most important of them.

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 robotic 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—each carrying a lithium-polymer pack, with the two sides’ stated capacity and plans for 2026 together exceeding ten million units a year—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, priced against a link budget that has been computed rather than assumed, for 1.4 per cent; the weapon itself for 0.4 per cent. A system that denies ground without discriminating costs 110 mWh per emplaced-day and one that looks first costs about eleven times that—and across the full stated uncertainty the multiple never falls below about six and a half—so **what the energy purchases is discrimination**. A mine is the sensorless limit of the same design: it performs no sensing 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 sensorless case. On the figures here, detectability costs approximately nothing to meet in the energy-spending case; self-neutralization comes cheaply from the architecture and self-destruction is also cheap; self-deactivation is in tension with the architecture rather than free; and discrimination—the obligation the law does not impose on mines, because no device without sensing could meet it—costs about 1.1 Wh 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. On a channel that carries a command to *stop*, jamming denies recall without preventing function, so it makes the system more dangerous rather than less; on the permission-to-continue channel the paper recommends at §25.3, the same jamming safes it within one authorisation period. **The direction in which a communications failure operates is set by the architecture and not by the adversary**, which is the finding this paper would put first. It follows 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 active ranging in the passive baseline**, which removes a continuous watt-level ranging load from both the emission and the power budget. That is narrower than removing lidar: one-time survey lidar, event-driven lidar, scheduled search and continuous search are four different architectures with four different energy forms and emission consequences, and §15.5 separates them. **The one-time survey improves the metric landmark set on which the interlock below depends and preserves the passive baseline as costed only if its energy is supplied externally or excluded by a stated system boundary, since a five-minute survey drawn from the emplacement battery is about 6.6 per cent of the daily budget once amortised (§15.5.2); event-driven lidar adds a confirmatory branch that must be priced; continuous active search abandons the premise of Part IV and is a different architecture, not a variant of this one.** “Military-grade” names no range: what a ranging sensor achieves is a function of wavelength, aperture, detector regime, reflectance, incidence angle, atmosphere and the false-alarm rate the operator tolerates, and the paper parameterises it rather than quoting a figure. 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 a **survey-integrity interlock**: a landmark set measured once at emplacement, a robust pose fit against it on each wake, an always-live inertial channel to catch the displacement at the moment it happens, and a release condition that expires when the two cannot agree that the platform still occupies its surveyed frame. Both channels are closed in the sense of §6.1—an inertial unit depends on no third party—and the pairing is used because they fail differently: the inertial channel sees displacement in darkness and fog where the visual channel cannot evaluate, and only the visual channel can *clear* an alarm, because clearing requires knowing the sector still covers the surveyed ground.

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 seventy 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, 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 also compiles what the public record—budget exhibits, the federal award record, granted patents and the open technical literature—contains for this capability class, under a rule that lets it say what is funded, what is on sale and what is physically calculable, and nothing about performance. What that record establishes can be said plainly, and what it does not establish should be said in the same breath: **the record shows that the component functions are separately public, funded or sold, and shows nothing whatever about the range, classification accuracy, reliability or integrated performance of the assembled system.** Subject to that, everything the specified object requires already exists; most of it is cheap, with the platform, the sensor and the effector each listed to the public at consumer prices; its components are offered for civilian sale in the United States, subject to the jurisdictional, export and end-use restrictions that attach to them; and in the systems now sold and fielded the safeguard is a layer that can be removed, so the device that does not discriminate is the simpler and cheaper one, and a safeguard that is to be more than a setting must be architecturally independent of the sensing it restrains, which is how Part V specifies it. What that record does not contain is which arrangements any party has built and how well they work, and the record marks that boundary itself; it is the reason a constraint drafted against function and against use holds where one drafted against named programmes or named technologies does not.

Against that record the paper sets a second one, assembled by function rather than by system. Waiting unattended with a live sensing system, matching an observed signature against a stored profile, engaging without a person designating the particular object, and re-tasking an emplaced munition from the rear are each separately dated to a fielded programme, the oldest to 1979 and none later than 2009; autonomous engagement of materiel targets in local defence is not merely fielded but carries an explicit exemption in national policy and is operated by at least thirty states. One finding follows from that record and it is the paper’s second contribution: **in every fielded system for which the sequence is documented, the autonomous capability came first and the human control was added afterwards as a mode**—a manufacturer describing its safeguard as a customer-requested retrofit to an auto-firing baseline, an air-defence system whose two legal characterisations are a switch position, and one government’s test authority recording that a fielded munition’s autonomous engagement capability was removed by policy. That last case is the useful one, because what can be removed by policy can be required to be absent by law. The paper also records what the text now before the treaty body asks: measures to intervene in the operation of such systems, to the extent possible or feasible, including to deactivate or neutralize them in a timely manner, the December 2025 text’s alternative of incorporating self-destruct, self-deactivation or self-neutralization mechanisms having been dropped. On the figures here the surviving requirement depends on a command arriving through an adversary who can deny it, and the dropped one costs approximately nothing and is what the same body has required of a device with no sensing since 1996.

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 four 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.

The public record on the capability class—published government budget exhibits, federal contract award records, granted patents and trade press—is a fourth class, compiled in Part VIII under a rule stricter than any applied to the three above: where the record is a budget exhibit, an award record, a patent or a trade-press report, the grammatical subject of every sentence drawn from it is a document, a table, a record or a date, never a capability; where it is the open technical literature, every proposition is attributed to a named source, no equation is reproduced, and no relation is applied to any particular system. §31 states that method and its limits.

Authorities have been resolved against a primary or authoritative record before citation wherever that record could be reached. Where a source is cited for a proposition without having been read in full, the text says so at the point of citation; where a document could not be reached at all, Part VIII records the fact as a limit of retrieval rather than as a finding. Field-employment, procurement and public-record claims are current to the individual dates stated in their citations; the public-record compilation in Part VIII was made on 22 August 2026. **The substantive source review for this version closed on 27 August 2026**, subject to the expressly identified post-session check at §51.

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.

The companion papers. Two published papers carry the legal analysis this paper declines to make and take their technical figures from this paper: *A Robot Left in a Field Is a Mine the Law Cannot See*, DOI 10.5281/zenodo.21971255, and *The Interval, Not the Autonomy*, DOI 10.5281/zenodo.21970777, both open access under CC BY 4.0. Where this paper refers to a companion paper, it refers to these. Both were published against v1.0.0 of this paper (version 1, 17 August 2026, DOIs 10.5281/zenodo.21971256 and 10.5281/zenodo.21970778 respectively); version 2 of each

(15 October 2026) adds a postscript on the report the Group of Governmental Experts adopted on 4 September 2026 and leaves the body unchanged, as Appendix A.7 records.

Note on the two legal companion papers and corrected figures

Two legal papers published before this engineering paper cite figures from earlier working versions of it.¹ Their central propositions survive, but several printed numbers and one term do not. **This paper is the controlling engineering source.**

At the fifty-second worked readiness interval, the corrected figures are **1.221 Wh per emplaced-day** in total, **1.112 Wh per emplaced-day** for the identification-and-release-processing line, **110 mWh per emplaced-day** without that line, and a ratio of about **11.1×**. The command-to-stop receive architecture costs **17.0 mWh per day** under the stated traffic assumption; the six-hour permission-to-continue variant costs **6.7 mWh per day**. The latter is **not self-deactivation** in the definition used by Amended Protocol II, which requires the irreversible exhaustion of an essential component. These corrections alter the numbers and terminology, not the propositions that identification-and-release processing dominates the worked budget, that recall or bounded safing is cheap relative to that line, or that an emplaced sensing system raises the mine-regime questions examined in the companion papers.

Two further points of scope belong here because both companions depend on them. **The thirty-day horizon is a stipulation** borrowed from the hard figures Amended Protocol II's Technical Annex fixes for the relevant anti-personnel-mine categories; it is not an operationally measured optimum and does not generalise to all remotely delivered mines. And **the public-record proposition is bounded**: the component functions are separately public, funded or sold, and the record establishes nothing about the range, classification accuracy, reliability or integrated performance of any assembled system.

The engineering paper states and derives the complete mathematical framework. Each legal companion incorporates only the factual inputs and propositions on which its legal argument depends, normally by cross-reference. The papers must remain consistent in architecture, definitions and numerical inputs, but need not reproduce identical equations, terminology or drafting.

One thing this note does **not** do. The authority non-expansion invariant at §15.5.4 is an engineering design constraint proposed by this paper. It is not an obligation found in Amended Protocol II, in the report the Group of Governmental Experts adopted on 4 September 2026, or in either companion. The draft of 3 September 2026 had carried the two nearest provisions—that a system's target and mission parameters not be materially modified unless a human being executed or authorised the modification, and that modifications to existing systems with a material impact on the operational functioning of their target identification, selection and engagement functions be reviewed (CCW/GGE.1/2026/CRP.1/Rev.1, paragraphs 37(d) and 38(b))—and the adopted text carries neither (CCW/GGE.1/2026/3, paragraphs 25 to 44; the review provision is taken up at §15.5.5). Even those provisions left open what the invariant addresses: a change of sensing modality alters no parameter and may precede any review, and it may nonetheless enlarge the region within which the unchanged parameters operate. The invariant governs the envelope, and it is not

¹Christopher I. V. Farmer, 'The Interval, Not the Autonomy' <https://doi.org/10.5281/zenodo.21970778>; and 'A Robot Left in a Field Is a Mine the Law Cannot See' <https://doi.org/10.5281/zenodo.21971256>. Both were published before this paper and cite figures from working versions of it. The propositions they rely on are set out and, where corrected, identified in the note above.

to be read into the mine-law companion’s transplant of agreed mine-law obligations. §51 treats it accordingly.

Part I—Introduction

1. Statement of the problem

Endurance, as published for fielded unmanned ground vehicles, is a drive-hour figure: it states how long a platform can move. Small platforms are rated at four to eight hours and mid-sized logistics types at eight to sixteen;² the stated ambition for the next generation is endurance measured in days rather than hours.³ The mission examined in this paper is holding ground, and its defining feature is that the platform does not move. 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 published figure therefore does not describe what an emplaced system can do, and no published figure does. The discipline has no metric for that case, which is why it has no design method for it either.

The nearest reported case marks the boundary of what the record shows. A remotely operated platform, cued by uncrewed aircraft and recovered daily by its operators, held a position over about six weeks.⁴ That is evidence of employment persistence: the platform was kept in the role for that period. It is not evidence of unattended or offline operation, because it was remotely operated and recovered daily. What a system that is neither operated nor recovered could do is not established by the case; it is the question this paper takes up.

This paper 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. Limitations of the current recommendation

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

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.*

⁴*Militarnyi*, ‘UGV With a Machine Gun Replaced Infantry and Held a Position for a Month and Half’ (*Militarnyi*, 15 November 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, reporting the NC13 company of the 3rd Assault Brigade and stating that every morning the operators took the robot to the position for combat duty and returned it in the evening; and Patrick Tucker, ‘A Ukrainian ground robot defended a position from Russian assault for six weeks’ (*Defense One*, 19 May 2026), reporting what appears to be the same deployment—a Droid TW 12.7, the same unit lineage, the same mission—with a duration of 45 days in the body and six weeks in the headline, an operator some ten kilometres distant, and uncrewed aircraft overhead supplying continuous surveillance and real-time target information. The two accounts are not reconciled in either source and differ on whether the platform remained in position at all. Both are Ukrainian reports of Ukrainian performance and are treated as claims by a party to the conflict. The reports describe a remotely operated platform, and they are not evidence of unattended or offline operation.

platforms which do not plug into those networks “will sit in storage yards.”⁵ 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.⁶ 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. Scope and approach

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.

⁵Jorge 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.

⁶ibid, reporting Starlink data rates to ground robots as low as 10 Mbit/s with consequent degradation of video quality.

Part VIII compiles what the public record—budget exhibits, the federal award record, granted patents and the open technical literature—contains for this capability class, under a statement of method stricter than the rest of the paper applies, and records where that record stops and what it says about its own scope.

Part IX decomposes the engagement chain into functions, dates each to its first fielding from the official and manufacturers’ record, and states what the text now before the treaty body would and would not reach. It is the paper’s second empirical contribution and it is the evidence for the plain statement at §41.

Part X states the boundary—the limits of closure—identifies the interval within the emplacement at which the object changes what it is doing, and concludes.

4. Contributions and prior work

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.

A note on how these claims are stated. Where this paper says a thing is *new*, it means that a search of the sources listed at §31 did not find it, not that it does not exist; no systematic prior-art protocol was run, and a reader who knows the field may well place any of them. The claims are ordered so that the ones the paper would defend hardest come first, and the entries in §4.2 marked load-bearing are the ones it would not withdraw without withdrawing the paper.

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.⁷ The recent work on emissions control standard operating procedures for ground formations is likewise established practice being written down, not new theory.⁸

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.⁹ 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.¹⁰

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.¹¹

⁷See Part IV §11 below for the doctrinal sources and their limits.

⁸US 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.

⁹See Part VI §17 for the unattended ground sensor literature.

¹⁰See Part III §8.

¹¹See Part VI §19.

Priority is not claimed for any of these.

4.2 What is offered as new

Nineteen claims are listed below. Seven are load-bearing; the other twelve are findings that follow from them. A long list dilutes the few that matter, so the seven are enumerated here, marked [load-bearing] where they appear, and are countable against that mark: the **two-axis separation** with its withdrawal test (first); the **architectural inversion** that makes the endurance metric achievable (fourth); **energy as the price of discrimination** (sixth); the **compilation of the public record** under a stated rule (fourteenth); the **recall coupling** (fifteenth); the **heartbeat inversion** (sixteenth); and the **denial test** that generalises the last two (seventeenth). To those seven a later version added an eighth candidate, the **modality-independent authority gate** (nineteenth), which is marked load-bearing but is the newest and least tested of them. Withdraw any of these and the paper goes with it. The remainder are consequences, corrections or observations: useful, individually checkable, and not the reason to read it.

First, the two-axis separation and the withdrawal test (Part II). [load-bearing] 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 §16). 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 §17). [load-bearing] 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 §18). 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 §18 states the arithmetic and the assumption it rests on.

Sixth, energy as the price of discrimination (Part VI §21). [load-bearing] The budget re-read as a price list rather than as a constraint, with a mine as the sensorless 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 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 §14). 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 §15.4 and Part VI §16.2). 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 §17.3 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 a survey-integrity interlock over a metric landmark set: **a metric landmark set permits estimation of a rigid pose, and because a low-residual fit can describe either the original pose or a displaced one, safety depends on bounding the estimated pose change rather than on the residual alone.** Both channels used—visual landmark registration and an inertial unit—are closed under the test at §6.1, and the condition must gate weapon release rather than silently re-derive a sector.

Ninth, the detection cascade as the source of the energy budget (Part VI §17.4). 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 §30). 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. Networked munitions have already broken that in part, and §30 records the programmes that did: an emplaced munition whose *arming state* can be commanded from the rear is not new and is not claimed here. What is offered as new is that a closed holding system at posture 1 is the first obstacle for which *target* and *obstacle effect*—the two components those systems leave fixed—are themselves mutable after emplacement, because the object carries a discriminating sensor whose sector and release criteria are parameters rather than a detonation state. This is 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 §27). That jamming a *command-to-stop* holding system makes it more dangerous rather than less—while the same jamming safes a permission-to-continue one—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 §20). 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.

Thirteenth, the function chain and the finding that autonomy is a mode (Part IX). The discussion of these systems proceeds system by system, which imports a name, a nationality and a procurement history into every argument about what a machine does. Decomposing the engagement chain into functions—waiting, classifying, selecting, accepting, re-tasking—dating each to its first fielding, and pricing them against Part VI, is offered as new. So is the finding it produces: in every fielded system in the record for which the sequence is documented, the autonomous capability came first and the human control was added afterwards as a mode, and one government has documented removing such a capability by policy. That is what makes the function separable, and a function that is separable is one a rule can reach.

Fourteenth, and stated here because it is the paper’s plainest result, what the public record establishes in ordinary words (Part VIII §41). [load-bearing] Everything the specified object requires already exists; most of it is cheap; its components are offered for civilian sale in the United States, subject to the jurisdictional, export and end-use restrictions that attach to them; and in the systems now sold and fielded the safeguard is a layer that can be removed. Each clause is separately sourced in Part VIII and each is the kind of proposition a delegation can check. The compilation, under the rule at §31, is offered as new; the individual facts are not.

Fifteenth, the coupling of recall timeliness to jam resistance (Part VII §25.2). [load-bearing] That for a fixed occupied bandwidth and a fixed authenticated payload the two are one quantity, $G_p = 10 \log_{10}(B/L) + 10 \log_{10} t_{\text{msg}}$, so that **in the idealised processing-gain model** every doubling of the time a recall command takes to arrive adds approximately 3.01 dB of resistance to the jammer trying to stop it—acquisition, coding and implementation losses making the tabulated times lower bounds and the realised gain smaller. The two relations composed are textbook; their composition, and the reading of paragraph 35(c) that follows from it, are offered as new.

Sixteenth, the heartbeat inversion (Part VII §25.3). [load-bearing] That requiring a periodic authenticated permission to continue, rather than listening for a command to stop, reverses the direction of every communications failure, bounds how long the platform can remain armed without authorisation, costs less than recall-on-command, and occupies the functional position of a safeguard that operates when nothing reaches the system. Dead-man arrangements are old and the concept is not claimed; what is offered as new is the treatment of its failure direction, its energy price and its legal significance as one integrated result.

Seventeenth, the denial test (Part VII §29). [load-bearing] That the companion question to Part II’s withdrawal test is *what does an adversary obtain by denying this channel?*, that the natural design of every channel in this specification fails it, and that the three denial-safe repairs together cost under one per cent of the budget. Offered as new, and as the general form of which the two preceding entries are instances.

Eighteenth, the energy-constrained detection range and the usable field envelope (Part VI §17.5–§17.6). That a fielded system possesses not the range of its longest-range component but the intersection of the regions in which it can cue, identify, verify its own emplacement geometry, remain authorised and physically reach; that the cue term is bounded by energy rather than by sensing, because sensitivity and false-alarm rate move together and the false alarms are what the budget spends on; and that the resulting independent-range fraction separates what a closed

platform can support from what it can only reach on another party's cue. Offered as a framework; the experiment at A.2 is what would make it a measurement.

Nineteenth, modality-independent authority and the modality materiality test (Part V §15.5). [load-bearing] That a change from passive visual recognition to active metric sensing can change detection, identification, pose and emission performance without changing the system's procurement name or autonomy label; that the engineering control must therefore sit downstream of every sensing path rather than inside one of them; and that an automatic change of modality may preserve or contract human-conferred authority but may not enlarge its ceiling. The materiality of a sensor change is stated as a measurable vector Δ_m rather than as a label. Offered as new, and testable unarmed at A.11, which is why it is marked load-bearing despite being the newest claim here.

Twentieth, four consequences of pricing an active ranging sensor rather than naming one (Part V §15.5.1.1–§15.5.1.4). That robust wide-area recovery of ground beneath foliage is, in the published record of this device class, a multi-look result obtained by combining many viewing geometries, so that a fixed horizontal emplacement cannot assume it and a vendor's capability claim cannot be transferred to one without measurement; that receiver and readout power remain substantial however small the optical pulse energy, which disposes of any inference from single-photon sensitivity to a microwatt-class ranging system and rules out the one cited operating point as a continuously powered tier, and which makes receiver power a function of scene activity and therefore of weather; that the cost of a sector sweep follows from the instantaneous field of view, so that the figure of merit for an emplaced application is receiver power divided by array format; and that moving to an eye-safer wavelength moves the detector to a material system with a materially higher dark count rate, whose cost lands in the cascade or in the dwell according to where the modality sits. **None of these is a claim to have invented a device or to have characterised military lidar.** They are system-level relations, each resting on a full text or an attributed manufacturer specification, each offered with the conditions of its source attached, and each accompanied by the measurement at A.11 that would test it.

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 §17 states, the claim is wrong. That is a bench question and §17 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 §18 gives the thresholds.

The **scavenging proposal** fails if the recoverable fraction of expended airframes is materially below the figure §20 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 the discrimination line at §21.2 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_1 and q_2 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_2 can be lowered while P_2 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 **negative finding on documented autonomy** failed, and it failed in the manner §4.3 said it would: by a case in which the decision path was recovered and read. §50 records the withdrawal and what replaced it. The **narrowed** version—that no autonomous selection of a *person* has been independently documented—fails on the same terms, and a reader should expect it to fail sooner rather than later.

The **recall pricing** failed on its own link budget, which is the more instructive of the two: the channel was priced before it was shown to close, and Appendix A.9 then showed it did not close at the stated range. §25.1 records the correction. What would falsify the corrected figure is a measured receiver sensitivity materially worse than -120 dBm at the stated rate, or a message length materially longer than the two hundred bits §25.2 assumes.

The **modality-independent authority gate** fails in either of two ways, and both are checkable. It fails as a *contribution* if every materially different sensor path is already forced through a common authority gate in the architectures that matter—that is, if the problem it solves does not exist in practice, which a survey of real release architectures would settle. And it fails as a *proposition* if a modality change does not in fact alter any component of Δ_m beyond validated bounds, in which case sensor substitution is a component change and needs no repeated review. A.11 tests the second directly and the first by construction, since a rig that cannot be made to issue an out-of-authority request through any path is a rig in which the gate is doing nothing.

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 **function-chain finding** fails if a fielded system can be shown whose human-control requirement was architectural from the outset and whose autonomous mode was added later, or if the systems cited are shown to have been mischaracterised by the sources that describe them. It is a claim about a documented sequence and it is defeated by a counter-example. It would also be weakened, though not defeated, by evidence that the sequence reflects the order in which capabilities are described in marketing rather than the order in which they were engineered.

The **plain statement at §41** fails clause by clause and each is separately checkable: if a component of the specified object is shown not to exist, if the commodity listings are shown to be unrepresentative of what is obtainable, if the vendors' statements about lawfulness are shown to be wrong, or if the safeguards in fielded systems are shown to be architecturally independent rather than modal. The fourth clause is the one this paper would most want tested, and Part IX §49 is the evidence offered for it.

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. §16 states what it should discriminate between.

Part II—The Two Axes: Closure and Emission Posture

5. Two senses of “offline”

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.¹² 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

¹²The 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.

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.

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 makes no intended emission, it cannot be located by direction-finding of its own transmissions; the qualification a specialist reader will supply, concerning unintended emission from the receiver itself, is made once at §12 and narrows the receiver architecture rather than the conclusion. 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

8. Services consumed by a representative control stack

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.

8.1 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.¹³ Nothing turns on the composite being any particular vendor's platform, and the audit is more useful for not being.

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.¹⁴ 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.¹⁵ A request to extend coverage over occupied Crimea during an operation in September 2022 was declined by the company's principal.¹⁶ 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.¹⁷ 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.¹⁸

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.

¹³Rivero (n 5), describing control by satellite link, line-of-sight radio, fibre-optic cable and aerial relay according to range, terrain and electronic warfare conditions.

¹⁴'Starlink 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 PBS NewsHour, the Belfer Center and Abels, cited at nn 16–18 below.

¹⁵*ibid.*

¹⁶PBS 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.

¹⁷Belfer 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.

¹⁸Joscha Abels, 'Private infrastructure in geopolitical conflicts: the case of Starlink and the war in Ukraine' (2024) 30 *European Journal of International Relations* 842.

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.¹⁹ 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.

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.²⁰

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.²¹

¹⁹Rivero (n 5).

²⁰ibid.

²¹Army 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

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 §8.1'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.

8.2 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.²² That is a working existence proof for persistent autonomous sensing at a scale nothing in the military 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.²³

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.

²²ASD 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-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.

²³Anduril, 'Sentry' <https://www.anduril.com/sentry> accessed 5 August 2026, describing the platform's function in combination with the vendor's command software.

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 §15.1. 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 X §52 states the boundary in those terms, and the comparison should be kept in view whenever this paper’s architecture is described as sufficient.

8.3 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 13) and the classes at §8.1. 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 nine 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.

9. Limits of the audit

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 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

10. Scope and limits of Part IV

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.

One taxonomic point belongs here rather than in Part V, because it governs the whole Part. Emission in this Part means intended radiated energy, and it includes optical emission. A ranging laser and a radio transmitter both radiate; both can be intercepted by an instrument matched to their band; and neither is categorically easier to detect than the other, because that comparison depends entirely on the interceptor. **This paper offers no ordering between an optical and a radio-frequency signature** and treats the active optical variants of §15.5.2 as members of the same taxonomy rather than as a separate category outside it.

And three different things are called detectability in this literature, which is a persistent source of confusion. *Intended emission* is what a transmitter or illuminator radiates on purpose, and is what this Part is about. *Receiver leakage* is what a receiver radiates unintentionally, and is the qualification §12 makes to the receive-only argument. *Post-conflict detectability* is a property required by Amended Protocol II’s Technical Annex so that an emplaced device can be found by ordinary mine-detection equipment after hostilities—eight grammes of iron or equivalent—and it has nothing whatever to do with either of the first two. **A design that satisfies the third is not thereby less covert in the tactical sense, and a design that is covert in the tactical sense does not thereby fail the third.** The Protocol’s requirement and a lidar pulse train are not the same subject.

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.

11. The governing asymmetry

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.²⁴ 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.²⁵ Naval practice reached the position earlier and holds it harder, because a submarine's survival has always depended on it.²⁶

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.²⁷ The asymmetry follows directly: the intended receiver applies the full processing gain of a matched filter, while the

²⁴The proposition is standard in emission control doctrine. For the ground-forces statement see US Army (n 8); for the naval statement see US Naval Institute (n 12).

²⁵US Army (n 8). 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.

²⁶US Naval Institute (n 12).

²⁷Don 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.

interceptor cannot, and is further penalised by having to accept a wide bandwidth and the noise that comes with it.²⁸

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.²⁹ 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.³⁰ That is the technical basis for the claim at §12 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.

12. Detection and location

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

²⁸US 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.

²⁹Harry 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.

³⁰*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.

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.

The claim that a system which only listens does not disclose itself needs one concession made at the outset, because a specialist reader will otherwise supply it. A superheterodyne receiver is not perfectly silent: it re-radiates its local oscillator, and that leakage has been used to locate receivers. The effect and its exploitation are documented in the peer-reviewed literature—C. Stagner, A. Conrad, C. Osterwise, D. G. Beetner and S. L. Grant, “A practical superheterodyne-receiver detector using stimulated emissions”, *IEEE Transactions on Instrumentation and Measurement* 60(4), 1461–1468 (2011); C. Stagner, “Detecting and locating electronic devices using their unintended electromagnetic emissions”, PhD dissertation, Missouri University of Science and Technology (2013); and, on detector performance, J. Manco-Vásquez and others, “Detection of radio receivers: an experimental evaluation approach”, University of Cantabria repository, handle 10902/7635—and the same technique is in routine commercial use in the devices sold to detect radar detectors. So the general proposition that direction-finding locates transmitters while a listener stays hidden is, stated without qualification, too strong: a listener with a local oscillator can be found.

That concession bears on the choice of receiver architecture, not on the conclusion, and it narrows the design space rather than widening it. A tuned-radio-frequency receiver has no local oscillator and therefore has nothing to leak, and the tuned-radio-frequency stage is precisely the architecture on which the paper's own lowest-power figure depends (§25.1). The qualification thus points to the design the paper already recommends: the leakage mechanism that betrays a superheterodyne listener is absent by construction from the receiver the paper specifies. The correct statement is the narrow one—a tuned-radio-frequency receiver is not detectable *by this means*—and not the broad claim that it is undetectable. Even so narrowed, the position is stronger than the unqualified claim it replaces, because it survives the objection a specialist would raise and lands on the architecture the rest of the paper is built around.

13. Comparison of emission postures

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 makes no intended emission, and this is the point the field has under-used. A receiver is not detectable by direction-finding of a transmission it does not make; the qualification about unintended emission from a receiver's own front end is made at §12, and the tuned-radio-frequency wake-up stage specified at §25.1 does not present it. The platform can be commanded, constrained and told to stop, while presenting no intended emission to distinguish it from 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.³¹

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.

14. Reporting thresholds and the transmission decision

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

³¹Rivero (n 5), reporting that operators struggle to maintain reliable links under intense jamming and that electronic fratricide degrades radio control.

none. The disclosure side is a function of the emission’s persistence and predictability, per §12, 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 §15.6 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 §17 proposes and §18 prices.

Part V—Specification of the Architecture

15. The holding platform

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. The thirty-day period is a stipulation, adopted from the horizon Amended Protocol II’s Technical Annex fixes for the self-destruction of remotely delivered *anti-personnel* mines (§21.1); it is not a figure derived from any operational report. Where a design choice depends on that case, the dependence is stated.

15.1 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.³² A study of seismic detection and tracking, using a three-component velocity transducer, reports a footstep trial with a 75 kg walker in open grass and a tracking trial in which the peak signal from a jogger on a 50 m radius sits well above the noise.³³ The physics sets the shape of the decline, and that study states it as a transmission loss with two terms: a geometric term falling as the logarithm of range, and an excess attenuation term proportional to the product of range and frequency, with an attenuation constant of about 30 dB per 100 m at 100 Hz. Two consequences follow from that form and both matter to the design. Geometric spreading acts differently on the two wave types and the quantity has to be named or the comparison is meaningless. **Body-wave amplitude falls as the inverse of range, which is 6 dB per doubling of distance; Rayleigh surface-wave amplitude falls as the inverse square root of range, which is 3 dB per doubling.** Energy falls as the square of each. The surface wave therefore reaches further by construction, and the same study reports that the Rayleigh amplitude is detectable at greater range than the body-wave envelope; it also apportions the energy of an impact at roughly two-thirds into the Rayleigh wave. And because the excess term scales with frequency as well as range, **detection range is frequency-dependent:** the low-frequency content of a footfall propagates furthest, which makes any single range figure a property of the ground and of the passband rather than of the sensor.³⁴ 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.³⁵

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.

³²Gökhan 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.

³³George Succi, Gervasio Prado, Robert Gampert, Torstein Pedersen and Hardave Dhaliwal (SenTech), ‘Problems in seismic detection and tracking’ in *Unattended Ground Sensor Technologies and Applications II* (SPIE 2000), read in full. The paper gives seismic transmission loss as a geometric term plus an excess attenuation proportional to range and frequency, with an attenuation constant of approximately 30 dB per 100 m at 100 Hz; states 6 dB per doubling of distance for body waves; reports the Rayleigh-wave detection range as greater than that for the body-wave envelope; and describes footstep trials with a 75 kg walker in an open grass field and a tracking trial at a 50 m radius. Earlier versions of this paper attributed to unnamed related work a footstep detection range of approximately 94 m and a stealthy-footfall figure of approximately 51.5 m. Those figures are not in this source, the related work was never named, and neither figure could be resolved to a document; both are withdrawn, and Appendix A.7 records the withdrawal.

³⁴ibid. 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.

³⁵Koç and Yegin (n 32).

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 §17.4 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 40 s per sweep, being the settle-and-capture figure of §17.3.]*

Sweep cadence	Energy per day	Over 30 days	Harvesting surface	Sweep energy as a multiple of the passive budget
Hourly	2.7 Wh	80 Wh	7.2 Wp	×2.2
Every 10 minutes	16.0 Wh	480 Wh	43 Wp	×13
Every 2 minutes	80 Wh	2.4 kWh	216 Wp	×66

An hourly sweep adds roughly twice the passive budget on top of it and remains within everything Part VI claims. A ten-minute cadence adds 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 §18.1, 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 §15.9, 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 §17.

Not continuously present in the baseline: active sensing. No radar and no continuous laser ranging. Both are emissions in the sense of Part IV and both are detectable at ranges greatly exceeding their useful ones. **Four arrangements have to be kept apart here and §15.5 does so:** a one-time ranged survey at emplacement, which the baseline admits and which §15.4 relies on; event-driven active confirmation after a passive escalation, which is a priced variant; scheduled

active search, which is a different mission posture; and continuous active search, which is a different architecture and is not costed in this paper’s budget. What is off by default is the third and fourth. Active sensing is otherwise available as a commanded mode where the tactical situation justifies the exposure.

15.2 The optical stage and Protocol IV

This paper specifies an optical sensing stage, and it should not pass over the treaty protocol that governs laser weapons; silence here, in a paper that is nowhere else silent, would read as an oversight before an audience more literate on that instrument than any other. The instrument is Protocol IV to the Convention on Certain Conventional Weapons, on Blinding Laser Weapons, of 13 October 1995.³⁶ Its prohibition is on laser weapons specifically designed, as their sole combat function or as one of their combat functions, to cause permanent blindness to unenhanced vision. It further requires parties to take all feasible precautions to avoid incidental permanent blindness in the employment of laser systems.

The prohibition in the first limb attaches to design purpose, and not to wavelength, to power, or to the possibility of harm. On the face of the instrument a sensing system does not fall within the prohibition by reason of being optical, and it would not be taken outside the prohibition by being weak. The paper makes no engineering claim about eye safety of any kind: it has measured nothing, and it neither states nor implies that any wavelength, power level or aperture is safe or unsafe. The point here is only about how the first limb is framed—it turns on the purpose for which a system is designed—and nothing in that framing depends on a safety figure the paper is not qualified to give.

The design consequence lies in the second limb. The obligation to take all feasible precautions against incidental permanent blindness attaches, by its own words, to “the employment of laser systems”. It is not engaged by an architecture merely because that architecture is optical: a passive electro-optical or thermal stage emits nothing and is not a laser system, and the baseline at §15.1 excludes laser ranging precisely. **Protocol IV does not supply an engineering safety classification, and the two must not be run together.** The instrument states a prohibition and a precaution obligation; it does not classify a device, set an accessible-emission limit or tell a designer what is safe. That is the province of the applicable laser-product standard, against which wavelength, accessible emission, aperture, dwell and scanning behaviour must be assessed for the actual design.³⁷ A system may satisfy the standard and still engage the Protocol’s second limb, and a system may fall outside the Protocol’s first limb and still be unsafe. §15.5 fits a ranging sensor and therefore raises both questions; this paper answers neither, and says so.

The obligation is engaged the moment a laser is installed or employed—the commanded range-finder that §15.4 admits at emplacement, a designator, or any illuminator a variant might add—and it is engaged whether or not the first limb is. A designer who adds one has entered the Protocol’s second limb and should know it; a designer who does not, has not. The Protocol further records at its third article that blinding as an incidental or collateral effect of the legitimate military employment of laser systems, including against optical equipment, is not covered by the first limb’s

³⁶Protocol on Blinding Laser Weapons (Protocol IV to the Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons), Vienna, 13 October 1995, arts 1 and 2.

³⁷IEC 60825-1, *Safety of laser products — Part 1: Equipment classification and requirements*. **Not consulted.** The standard is named here because it, and not Protocol IV, is the instrument that classifies a laser product, and because a designer must assess the actual design against the edition and amendments then in force. This paper has not obtained the standard, states no classification, and makes no safety claim of any kind.

prohibition. This paper takes no position on whether any arrangement it describes satisfies either limb of the Protocol; it states what the instrument says and what follows for a designer, and stops there. It is worth recording that the companion legal papers cite this Protocol twice, and that this paper had not cited it at all until this passage.

15.3 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 §15.1 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.

15.4 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 is relieved of that continuous burden by a survey performed once. What exactly it is relieved of is set out at the end of this subsection, because an earlier version of this paper claimed more than the argument supports.

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 §15.1 excludes from the baseline 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 §17.3 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 §15.6 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, and stating it precisely matters because the loose version of it is wrong. At emplacement the platform surveys the static scene: a set of fixed features—a wall corner, a treeline break, a culvert mouth, a burnt-out hull—recorded as a rigid three-dimensional landmark set in the platform’s own frame. Thereafter, on each tier-three wake, it re-observes those landmarks and asks whether a single rigid change of its own pose accounts for what it sees.

Three properties of that test decide whether it works, and the first two rule out the intuitive version of it.

Bearings alone are not enough, because a translation and a change in the world produce the same first-order symptom. Under a pure rotation of the platform every bearing shifts by the same angle and the angles *between* features are preserved. Under a translation they are not: near features shift more than far ones, which is parallax, and the included angles change by amounts that depend on range. So “every bearing has shifted, therefore the platform moved” is not a sufficient rule—it is only sufficient for rotation, which is the least likely displacement and the least consequential.

And range is not free from one passive view. A single monocular image supplies bearings, not distances, so the pairwise ranges the test needs must be established once, at emplacement, by a method that measures them. Three are available and the choice is a design decision rather than a difficulty: a stereo baseline from a sensor pair; a short commanded translation of the platform itself, giving structure from known motion; or a single ranged sweep with a laser range-finder at the moment of emplacement. **The third is worth stating because it looks like a contradiction and is not.** §15.1 excludes active ranging from the *baseline* on emission grounds, and Part IV’s argument is about *persistent* emission: a single sweep at the moment a human or a carrier robot is already standing at the site adds one brief emission to a moment that is already the loudest in the platform’s life, and adds nothing to the thirty days that follow. The emission budget at Part IV is a budget over the emplacement, not over the instant of emplacement.

With a metric landmark set in hand the test is a standard and well-posed one: solve for the platform’s pose from the observed landmark bearings, with robust outlier rejection, and read the answer off the residuals.

What the fit shows	Reading
A single pose explains nearly all landmarks, and it is not the emplacement pose	The platform moved , by that pose
No pose explains a minority of landmarks while the rest sit at the emplacement pose	The world changed in the inconsistent subset
No pose explains the majority	Ambiguous: revert to safe

The third row is the one that makes the method safe rather than merely clever. A test that must always return one of two answers will return the wrong one under fog, snow, a burnt scene or a

demolished wall; a test permitted to return *ambiguous*, and required to safe the platform when it does, fails in the direction §26 argues for.

Stated as an interlock, because that is what it is. The three readings above are a decision procedure, and a decision procedure that gates weapon release should be written as a condition rather than described. Let the fitted pose be $(\hat{t}, \hat{R})$ and the emplacement pose (t_0, R_0) ; let m be the inlier count, ρ the fit residual, D_{IMU}^u an uncleared inertial displacement indication, and $f_V(t)$ a freshness predicate recording that the visual or metric check has been evaluated recently enough to be relied on. The platform holds **survey integrity** at time t when

$$V(t) = (m \geq m^*) \wedge (\rho \leq \rho^*) \wedge (\|\hat{t} - t_0\| \leq \delta_t) \wedge (d_{\text{SO}(3)}(\hat{R}, R_0) \leq \delta_R) \wedge \neg D_{\text{IMU}}^u \wedge f_V(t)$$

Translation and rotation are separate terms because they have different units and different consequences, and an earlier version of this paper wrongly combined them in one norm. It also wrongly said that a displacement small enough to pass the pose bound would be caught by the residual. It would not. **A rigidly displaced sensor observing the same rigid landmark set fits it perfectly:** the residual measures how well a pose explains the observation, not which pose it is. The residual and inlier count validate the *fit*; the estimated pose change is what detects the *displacement*. Conflating the two is the commonest way this class of test is got wrong, and stating them as separate clauses is what prevents it.

The thresholds are set by the boundary, not by the sensor. δ_t and δ_R should be derived from how far the authorised boundary may be permitted to move, not chosen for the convenience of the estimator. For a small rotation and a limiting effect distance R_{eff} , the displacement of a boundary point is bounded to first order by

$$\Delta_b \lesssim \|\Delta t\| + R_{\text{eff}} d_{\text{SO}(3)}(R, R_0), \quad \Delta_b \leq b^*$$

so a metre of tolerated boundary movement at 30 m is bought with roughly a metre of translation *or* about two degrees of rotation, and the rotation term dominates at any useful range. **The approximation holds for small rotation and is a design bound rather than a measured error law;** A.5 measures the real one.

Three properties of the condition are the reason for writing it this way rather than leaving it as a description.

It is a conjunction, so every clause is a veto and none is a vote. A rotation that moves the sector is caught by δ_R whatever the residual says; a scene change that destroys the landmark set is caught by m ; a displacement in darkness or fog that leaves no visual evidence at all is caught by D_{IMU}^u ; and a check that has not been run recently enough to mean anything is caught by f_V . The failure of any clause—including the failure to *evaluate* a clause, which is what an obscured scene produces—is a failure of V .

It converts an observation into an assurance mechanism, and an assurance mechanism can be tested. The experiment is at Appendix A.5 and needs no weapon: an instrumented, unarmed platform, a surveyed landmark set, and a schedule of controlled displacements and controlled scene alterations. What is measured is the pair of error rates—the rate at which V survives a displacement that should have expired it, and the rate at which V expires without one—against displacement magnitude, landmark count, and the visibility conditions under which the check runs. **The first of those two rates is the safety-relevant one and it is the one a specification should bound.**

And it is the boundary of the envelope at §17.6. C_{pose} in that intersection is the region over which V can be evaluated and holds: a sector arc which the surveyed landmark set does not cover is a region in which the platform cannot establish that it still occupies its own frame, whatever its sensors can see there. That is a constraint on where an emplacement may be authorised to act which does not appear in any specification this paper has examined, and it is a property of the *site* rather than of the machine.

Inertial sensing belongs in this design, and an earlier version of this paper excluded it on a mistaken ground. That version preferred visual self-survey over an inertial unit on the basis that it “passes the withdrawal test of §6.1, which is not true of any of the alternatives”. That is true of satellite positioning, of a beacon and of a survey party. **It is not true of an inertial measurement unit, which depends on no third party at all and passes the withdrawal test as cleanly as anything in this specification.** The error was to lump a closed sensor with open services because both were “alternatives”, and it is corrected here because closure is this paper’s own central distinction and misapplying it in the paper’s own terms is worse than not raising it.

The real objection to inertial sensing is drift, and drift disqualifies it only from the job it was never needed for. An inertial unit cannot tell the platform where it is after thirty days. It can tell it, immediately and for a few microwatts, that *something just moved it*—a shock, a tilt beyond threshold, a sustained acceleration. That is a detection problem, not a navigation problem, and it is the cheap always-on tier of exactly the architecture this paper argues for everywhere else:

Tier	Function	Cost
Inertial, always live	Detect that a displacement event occurred; safe the platform at once	microwatts, continuous
Visual landmark fit, on wake	Quantify the displacement, or clear the alarm and re-authorise	runs on imagery tier three has already captured

The two answer different questions and the pairing is strictly better than either alone. The inertial channel catches the displacement at the moment it happens, including at night, in fog, and when the scene is obscured—cases in which the visual test returns *ambiguous* and would otherwise safe the platform for a benign reason. The visual channel catches the displacements the inertial channel misses—slow subsidence, freeze–thaw heave, a slow drag—and is the only one of the two that can *clear* an alarm, because clearing requires knowing the sector still covers the surveyed ground and an inertial unit cannot know that. Neither channel depends on anything outside the platform.

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

§19. Part VII §27 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 visual channel. It requires the scene to contain static features to survey, and on genuinely featureless ground there is nothing to hold a pattern. There the inertial channel carries the function alone: it still detects that a displacement occurred and still safes the platform, and what is lost is the ability to *clear* the alarm without a person. That is the correct place for the limitation to fall.

The dividend is real but narrower than an earlier version claimed. That version said a static platform “obtains the same discrimination from a survey performed once” as a moving one obtains from lidar and simultaneous localisation and mapping. It does not. What stasis removes is the need to estimate ego-motion continuously in order to separate the platform’s own motion from the world’s, and with it the continuous active ranging that estimate normally requires. What it does not supply is metric range to a newly detected object: background subtraction establishes that something in the scene has changed, not how far away it is or what shape it has. Range to a target, where the release criterion needs it, must come from the same one-time survey geometry, from a stereo baseline, or from the known size of a recognised object class—and the last of those is an inference, not a measurement, and should be treated as one. The saving is large and it should be claimed at its true size: **stasis removes a continuous watt-level ranging load and a continuous estimation load, and it buys change detection rather than scene understanding.**

15.5 Active metric sensing, modality substitution and authority non-expansion

The baseline at §15.1 is passive, and Part IV’s emission argument is the reason. But a passive baseline is a choice about *this* design, and the choice is reversible by anyone who buys the same platform and fits a different sensor. This subsection sets out what an active metric sensor would change, what it would not, and the one property the specification must hold invariant across the change. **It is the part of the specification most likely to be tested by events, because substituting a sensor is the cheapest way to alter what a fielded system can do while leaving its name, its label and its procurement description untouched.**

15.5.1 *What active 3D ranging offers, parameterised rather than asserted*

“Military-grade lidar” is an informal umbrella and not a specification. It states no range, because range is not a property of a class of device: it is a function of transmitted energy, aperture, detector regime, target reflectance, incidence angle, atmospheric transmission and the false-alarm rate the operator will tolerate. **A useful description states the parameters**—wavelength; pulse or modulation scheme; aperture; detector regime; angular resolution; range precision; field of view; repetition or scan rate; processing method; size, weight and power; and emission and safety constraints—and this paper uses that list wherever it discusses the sensor, in place of a single number.

Within that framing, the open literature now examined in full establishes what the architecture is, what it costs and where its limits sit. The foundational demonstration is a compact Geiger-mode three-dimensional imaging laser radar whose published design parameters are a 532 nm passively Q-switched microchip laser at five to ten kilohertz, a 32 by 32 array of Geiger-mode avalanche photodiodes integrated with timing circuitry above 500 MHz, a 7.5 cm receive aperture at f/4.0, a projected pixel instantaneous field of view of 333 microradians, a total instantaneous field of view of 10.3 by 10.3 milliradians, and 15 cm range resolution.³⁸ The camera-level integration of such focal planes is separately published, with continuous streaming at frame rates of the order of two hundred kilohertz, range gates spanning four nanoseconds to ten microseconds, and automated bias control that holds photon detection efficiency constant against ambient temperature change.³⁹ The detector array and readout technology has its own twenty-year published record,^{40 41} and the underlying single-photon detection physics—photon detection efficiency, dark count rate, afterpulsing and timing jitter—is the body of work Part VIII already relies on.⁴² Part VIII §32 records three federal awards for photon-counting and Geiger-mode ladar receivers, and this paper draws from them what

³⁸R M Marino, T Stephens, R E Hatch, J L McLaughlin, J G Mooney, M E O'Brien, G S Rowe, J S Adams, L Skelly, R C Knowlton, S E Forman and W R Davis, 'A compact 3D imaging laser radar system using Geiger-mode APD arrays: system and measurements' in *Laser Radar Technology and Applications VIII* (2003) Proc SPIE 5086 <https://doi.org/10.1117/12.501581>. **Full text examined.** A copy was supplied to the author; the publisher's automated-access control was not circumvented, and Appendix A.7 records the change in evidence class. **The PDF's original text layer lacks a reliable Unicode mapping and corrupts alphanumeric extraction. Numerical values were recovered by forced-raster OCR of the rendered pages and verified visually against the page image; none is taken from unverified text extraction.** Design parameters are from Table 1 at page 2 of the proceedings pagination: 532 nm laser, 5 to 10 kHz; 32 × 32 Geiger-mode APD array with CMOS timing above 500 MHz; 7.5 cm receive aperture, 30.0 cm effective focal length, f/4.0; projected 100 μm pixel pitch of 333 μrad; field of view 10.3 mrad × 10.3 mrad; range resolution 15 cm. Transmitter characteristics are from Table 3 at page 4, whose column is headed *Measured Value*: 532 nm, pulse length about 700 ps, **energy per pulse 33 μJ at 532 nm**, pulse repetition rate 5 to 10 kHz. The wide-angle composite of about 8 mrad in elevation by about 160 mrad in azimuth, assembled from multiple step-stare collections positioned by pan-tilt encoder readouts, is at page 5. **The system is a 2003 laboratory and field testbed and is not a fielded military design point**; nothing here is offered as a specification for any current system.

³⁹M Entwistle, M A Itzler, J Chen, M Owens, K Patel, X Jiang, K Slomkowski and S Rangwala (Princeton Lightwave Inc), 'Geiger-mode APD camera system for single-photon 3-D LADAR imaging' in *Advanced Photon Counting Techniques VI* (2012) Proc SPIE 8375, 83750D <https://doi.org/10.1117/12.921004>. **Full text examined.** A copy was supplied to the author; the publisher's automated-access control was not circumvented. Cited for the camera-level integration of Geiger-mode focal-plane arrays: 32 × 32 InGaAsP GmAPD arrays optimised for 1.06 μm or 1.55 μm, frame rates in excess of 250 kHz for 0.4 μs range gates, continuous streaming at about 186 kHz with 2 μs range gates and 13-bit time stamps, range gates spanning 4 ns to 10 μs, and automated bias control holding photon detection efficiency constant against ambient temperature change. **This note replaces, as the load-bearing citation for camera architecture, M A Itzler and others, 'Geiger-mode APD single-photon cameras for 3D laser radar imaging' in 2014 IEEE Aerospace Conference 1–12** <https://doi.org/10.1109/AERO.2014.6836476>, whose full text was not obtained and which is an **IEEE rather than an SPIE record**; Appendix A.7 records the substitution.

⁴⁰B F Aull, E K Duerr, J P Frechette, K A McIntosh, D R Schuette and R D Younger, 'Large-format Geiger-mode avalanche photodiode arrays and readout circuits' (2018) 24(2) *IEEE Journal of Selected Topics in Quantum Electronics*, article 3800510 <https://doi.org/10.1109/JSTQE.2017.2736440>. **Full text examined.** A twenty-year MIT Lincoln Laboratory record of silicon and InP Geiger-mode avalanche photodiode arrays and the CMOS readout circuits that count or time-stamp single-photon events, from 4 × 4 arrays to 256 × 256.

⁴¹B F Aull, 'Geiger-mode avalanche photodiode arrays integrated to all-digital CMOS circuits' (2016) 16(4) *Sensors* 495 <https://doi.org/10.3390/s16040495>. **Full text examined.** A review of the same programme, cited for the detector and readout record and for afterpulsing and crosstalk as the principal disadvantages of the device class.

⁴²S Cova, M Ghioni, A Lacaita, C Samori and F Zappa, 'Avalanche photodiodes and quenching circuits for single-photon detection' (1996) 35(12) *Applied Optics* 1956–1976 <https://doi.org/10.1364/AO.35.001956>. **Full text examined.** Cited for the characterisation of single-photon avalanche diodes by photon detection efficiency, dark count rate, afterpulsing and timing jitter, and for the dependence of afterpulsing on quenching-circuit design. The

it draws from every award record: that the function is funded and contracted, and nothing about delivered performance.

One point of attribution belongs here, because the paper’s own cascade is in the literature. The two-stage arrangement this paper specifies—a wide-field passive sensor that cues, and a narrow-field active ranging sensor that confirms—is not new. It is the seeker architecture described in the Geiger-mode lidar literature from 2002, where the stated goal was a system “that featured both a passive infrared imaging sensor with a wide field of view and a high-resolution lidar sensor with a narrow field of view”, and where the applications named alongside interceptor seekers are combat identification, navigation of autonomous robots and vehicles, surveillance and topography.⁴³ **The paper claims no novelty for the cascade.** What it claims is the accounting: what the second stage costs in energy, what it buys in envelope, and what must be true of the control that sits downstream of both.

What such a sensor can *potentially* supply, and what the passive baseline cannot, is a short and specific list:

- **direct metric range**, rather than range inferred from a monocular image and an assumed object size;
- **three-dimensional point clouds and object dimensions**, from which geometry rather than appearance is measured;
- **spatial relationships and registration** against ground and landmarks, which is exactly what §15.4’s survey depends on;
- **motion estimation** from successive observations of the same structure;
- **geometric confirmation or rejection** of a hypothesis a camera has already formed;
- **operation independent of scene illumination**, subject to atmospheric transmission, target reflectance, incidence angle and occlusion;
- **under one specific and demanding condition, partial recovery of ground beneath foliage**—but the condition is the reason this paper does not claim the capability for its own design case, and §15.5.1.1 sets it out;
- and, in consequence, **a materially different detection and identification envelope** from passive electro-optical or thermal sensing.

And the limitations belong in the same list rather than in a later caveat, because the first two are decisive. Four of them are then developed at §15.5.1.1 to §15.5.1.4, because each turned out to carry a consequence for the energy budget, the emission argument or the legal analysis that could not be seen from the summary.

same source is relied on in the Part VIII compilation, where it was previously described as cited and not read in full; that description is corrected.

⁴³M A Albota, B F Aull, D G Fouche, R M Heinrichs, D G Kocher, R M Marino, J G Mooney, N R Newbury, M E O’Brien, B E Player, B C Willard and J J Zayhowski, ‘Three-dimensional imaging laser radars with Geiger-mode avalanche photodiode arrays’ (2002) 13(2) *Lincoln Laboratory Journal* 351–370. **Full text examined.** No DOI is assigned to this journal issue. Cited for the applications the authors name—advanced interceptor seekers, combat identification, navigation of autonomous robots and vehicles, surveillance and topography—and for the stated seeker architecture of “a passive infrared imaging sensor with a wide field of view and a high-resolution lidar sensor with a narrow field of view”, which is the cue-then-confirm cascade this paper specifies. **The paper is a 2002 laboratory and brassboard record; nothing in it establishes fielded performance.**

Because it emits, active ranging changes the emission posture and the energy budget; it does not by itself change the system’s closure. A ranging sensor that transmits is detectable in the sense §11 describes, and it draws power continuously if it searches continuously. Both of those are consequences on Part II’s *emission* axis and on the budget of Part VI. **The dependency axis is untouched:** a lidar that answers to nothing outside the platform is exactly as closed as a camera that answers to nothing outside the platform, and a modality that required a third-party service would fail §6.1 whether or not it emitted. The paper’s emission argument and its energy argument are premised on a platform that does not transmit to see; its closure argument is not.

Its adverse-weather behaviour is a serious limitation, and the evidence must be stated with its qualifiers or it will be misused. In Bijelic and others’ vehicle-detection benchmark, lidar-only average precision on the easy split falls from **73.46 in clear conditions to 28.98 in dense fog**, while the image-only detector’s reported AP remains approximately stable across the same splits. **The authors caution that differing scene complexity prevents treating this as a controlled modality ratio**—their own words are that “assessment metrics can increase simultaneously as scene complexity changes between the weather splits”, since fewer or less-occluded vehicles in bad weather can raise a score for reasons unconnected with sensing. The result therefore establishes **vulnerability in the tested lidar pipeline, not universal immunity of RGB sensing**, and the same table shows lidar-only AP at 67.50 in snow, so fog and snow are not interchangeable even within that benchmark. The same authors state, in introducing the problem rather than as a result of their own measurement, that in foggy or snowy conditions “state-of-the-art pulsed lidar systems are restricted to less than 20 m range due to backscatter”, and the qualifier belongs with the sentence: it is a statement about automotive-class pulsed systems, offered as motivation for the benchmark.⁴⁴ Rasshofer and others independently demonstrate the physical mechanism, reporting unwanted snow-generated target detections of **up to $2.5 \text{ s}^{-1} \text{ m}^{-2}$ within the near field—the first five metres—of the particular coaxial scanning automotive laser radar tested**, together with a maximum detection range against a 7 per cent reflectivity target falling from about 30 m in clear air to about 5 m at 20 mm/h of rain, and dry-snow attenuation of about 15 dB/km per mm/h.⁴⁵ **That raw-return density must not be equated**

⁴⁴M Bijelic, T Gruber, F Mannan, F Kraus, W Ritter, K Dietmayer and F Heide, ‘Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather’ in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 11679–11689 <https://doi.org/10.1109/CVPR42600.2020.01170>. **Full text examined** in the authors’ arXiv:1902.08913v3 of 30 June 2020, whose substantive content corresponds to the version of record; the DOI of record is cited. The figures quoted are the reported lidar-only average precision on the easy split: 73.46 in clear conditions, 67.50 in snow and 28.98 in dense fog, with the image-only detector approximately stable across the same splits. **The authors’ own caution is quoted with the figures and must travel with them:** assessment metrics can move because scene complexity differs between the weather splits, so the comparison is not a controlled modality ratio. The setting is automotive and the sensors are automotive-class; **the transfer to a military sensor is a transfer of the physical mechanism and not of the numbers**, and it is offered as such.

⁴⁵R H Rasshofer, M Spies and H Spies, ‘Influences of weather phenomena on automotive laser radar systems’ (2011) 9 *Advances in Radio Science* 49–60 <https://doi.org/10.5194/ars-9-49-2011>. **Full text examined.** At 905 nm the paper reports maximum detection range against a 7 per cent reflectivity target falling from about 30 m in clear air to about 5 m at 20 mm/h of rain, and dry snow attenuation of about 15 dB/km per mm/h against about 2 dB/km per mm/h for wet snow. In a coaxial configuration in heavy snow it reports that the transmit power density and the reflectivity of snowflakes generate “a false alarm rate of up to $2.5 \text{ s}^{-1} \text{ m}^{-2}$ ” within the first five metres in front of the sensor. **That quantity is a raw unwanted-target-detection density for the particular coaxial scanning automotive laser radar tested, in its near field, and is not a system-level false-alarm rate, a tier-one trip rate, a daily escalation count, or a property of lidar in general.** The physical mechanism transfers to other active optical systems; the numerical value does not transfer without testing the new wavelength, optics, scanning geometry, filtering, field of view and environment. §15.5.6 states the transfer relation and A.11 measures its coefficient.

with a system-level false-alarm rate; it is a count of unwanted detections at the detector, and its operational effect depends on filtering and on where the active modality enters the cascade. §15.5.6 sets out that dependence.

And two manufacturers publish a cold-weather power penalty, which bears directly on Part VI. One commercial sensor rated at 6.5 W average states that peak power “may reach 14 W” when self-heating engages between -20°C and 0°C ; another’s manual specifies a supply of at least 20 W and states that “if operating at cold conditions the peak power consumption is between 22 W and 28 W”, 28 W at -40°C , recommending not less than 30 W outdoors.^{46 47} **Roughly a doubling, in the season in which §18.1’s harvesting yield is at its annual minimum.** The active-modality term is therefore largest exactly when the energy available to pay for it is smallest, which is a coupling the passive baseline does not have. It is not the only path by which weather reaches the power budget: §15.5.1.2 sets out a second, through the dependence of receiver dissipation on the density of returns the weather itself generates.

A point cloud is geometry, not identity. Learned inference on point clouds is mature and efficient—end-to-end detectors operating directly on raw clouds without hand-crafted features are established,⁴⁸ and pillar-encoded architectures run at tens of hertz on commodity hardware⁴⁹—but what they produce is a class label from a trained model, subject to the same training data, threshold and distribution questions §21.2 raises about the passive tier. Multimodal benchmarks exist precisely because modality performance has to be compared on shared ground truth rather than asserted.^{50 51}

A longer sensing range is not an authority range, and §17.6 is where that is made precise.

⁴⁶Livox, ‘Mid-360 specifications’ <https://www.livoxtech.com/mid-360/specs> accessed 27 August 2026, stating power consumption of “6.5 W (average)” and that between -20°C and 0°C the unit “will automatically enter self-heating mode, where the peak power may reach 14 W”. **Manufacturer specification; not an independent measurement, not peer-reviewed, and not a military design point.**

⁴⁷Ouster, *OS2 hardware user manual for Rev7 sensors* (14 October 2024) <https://data.ouster.io/downloads/hardware-user-manual/hardware-user-manual-rev7-os2.pdf> accessed 27 August 2026, stating that the power supply must be “at least 20 W”, that “if operating at cold conditions the peak power consumption is between 22 W and 28 W” with 28 W at -40°C , and recommending “use of power no less than 30W if using in outdoor conditions”. **Manufacturer specification; not an independent measurement, not peer-reviewed, and not a military design point.**

⁴⁸Y Zhou and O Tuzel, ‘VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection’ in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 4490–4499 <https://doi.org/10.1109/CVPR.2018.00472>. **Abstract and method read; full text not read.** Cited for the proposition that detection can be learned end-to-end from raw point clouds without hand-crafted features.

⁴⁹A H Lang and others, ‘PointPillars: Fast Encoders for Object Detection From Point Clouds’ in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 12689–12697 <https://doi.org/10.1109/CVPR.2019.01298>. **Abstract and reported results read; full text not read.** Cited for the reported inference rates of 62 Hz, and 105 Hz for a faster variant, which is the basis for the statement that point-cloud inference is not computationally exotic.

⁵⁰H Caesar and others, ‘nuScenes: A Multimodal Dataset for Autonomous Driving’ in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 11618–11628 <https://doi.org/10.1109/CVPR42600.2020.01164>. **Abstract and dataset description read; full text not read.** Cited for the existence of a documented multimodal comparison setting: six cameras, five radars and one lidar with full 360-degree coverage over 1,000 annotated scenes.

⁵¹A Geiger, P Lenz and R Urtasun, ‘Are we ready for autonomous driving? The KITTI vision benchmark suite’ in *2012 IEEE Conference on Computer Vision and Pattern Recognition* 3354–3361 <https://doi.org/10.1109/CVPR.2012.6248074>. **Bibliographic record verified; full text not read.** Cited for the proposition that modality performance is compared on shared ground truth rather than asserted.

And Protocol IV is engaged, but engineering safety is a separate question. §15.2 states the treaty position: the precaution obligation attaches to the employment of laser systems, and fitting a ranging laser engages it where the passive baseline does not. Protocol IV does not, however, supply an engineering classification. Accessible emission, wavelength, aperture, dwell and scanning behaviour must be assessed against the applicable laser-product standard for the actual design,⁵² and eye safety must not be inferred from wavelength, from purpose, or from the fact that the instrument is a sensor rather than a weapon. **§15.5.1.4 returns to this and finds that the wavelength choice is not only a safety decision: it moves the detector to a different material system, and the cost lands in this paper’s own false-alarm term.**

15.5.1.1 Foliage recovery is a multi-look result at wide area, and one fixed viewpoint cannot assume it

The open record on seeing beneath vegetation with Geiger-mode lidar is more specific than the phrase “foliage penetration” suggests, and the specificity matters for this design. **Returns from beneath a canopy arrive principally through gaps in it.** What angular diversity buys is a higher probability that some line of sight to a given patch of ground is unoccluded, and a means of filling the volumes that remain shadowed from any one position. The peer-reviewed statement of the mechanism is explicit: the first system of its kind is described as “the first of a series of airborne lidar systems that could perform foliage penetration by fusing multiple 3D images taken from different look angles”.⁵³ Current commercial practice is the industrial form of the same idea—a polygon of ground scanned many tens of times from many tens of different viewing geometries, the scans then aligned and processed together to recover surfaces whose returns were too weak to detect in any one of them—and the operator chooses high angular diversity for foliated scenes and a single pass for open ground.^{54 55}

The foundational field experiment shows both halves of the position, and the paper states both. From a single sensor viewing direction the size and shape of objects under the trees *could* be measured; the shadow of the in-range foliage nevertheless appeared as holes in the scene, and the authors’ stated next step was algorithms “to accumulate the 3D data from a variety of sensor viewing directions with the motivation of filling-in some of the volumes that are shadowed or heavily obscured from any single viewing direction”. That experiment was conducted from an elevator on a 300-foot tower, looking down at the foliated ground at its base. The same paper records that

⁵²IEC 60825-1 (n 37). The standard, and not Protocol IV, is what classifies a laser product; this paper has not obtained it and states no classification.

⁵³Aull and others (n 40), describing the Jigsaw system.

⁵⁴D Fried, T Mangum and G Stamnes (3DEO Inc), ‘2024 Annual GeoChallenge: showcasing innovation with Geiger-mode lidar’ (2025) 91(9) *Photogrammetric Engineering & Remote Sensing* 535–539. **Full text examined; industry article, not peer-reviewed research, and no DOI.** Cited for the operational description of how foliage penetration is obtained: high angular diversity “crucial in highly foliated scenes to ‘poke through’ the foliage gaps to the ground layer”; a given polygon “scanned many tens of times from many tens of different viewing geometries”; a delivered point cloud defined as one “created after combining many scans together”; and the operator’s choice of high angular diversity for foliage and a single scan per polygon for unforested wide-area mapping. The article also states that a point is typically established from 5 to 10 detections arising from 15 to 30 received photons, against approximately 300 for a linear-mode system. **Manufacturer-authored; treated as a description of the method, not as independent validation of performance.**

⁵⁵D Fried and K Reichel-Vischi (3DEO Inc), ‘Next-generation Geiger-mode lidar systems’ *LiDAR Magazine* (30 December 2024) <https://lidarmag.com/2024/12/30/next-generation-geiger-mode-lidar-systems/> accessed 27 August 2026. The same authors’ fuller statement of the same method, including that the scans “are aligned, combined and then processed together in order to find surfaces whose lidar echoes were too weak” to be detected otherwise, and that collections use several flight lines at different headings. **Trade publication, manufacturer-authored; cited for the description of method only.**

for imaging through cover “the sensor actually works better with a relatively low probability of detecting the in-range obscurant or foliage”, arranged by reducing the laser power, apodising the receive aperture with an iris, or electronically reducing the sensitivity of the detector.⁵⁶

So the honest statement is narrower than a prohibition and wider than a capability claim. A single emplaced viewpoint may recover some under-canopy geometry through available gaps, as that experiment demonstrates, but it retains occluded volumes. **The robust wide-area recovery reported for the airborne class depends on angular diversity and multi-look fusion, and it cannot be transferred to a horizontal fixed emplacement without measurement.** Movement, retrospective combination and a downward look are what produce the published airborne result; the sources do not establish that they are absolute prerequisites for any under-canopy return at all, and this paper does not assert that they are. What a fixed emplacement *can* accumulate is time from one position, and foliage moves: a gap that is closed in still air may open in wind, and successive observations from one place are not nothing. **They are also not controlled angular diversity**, and the difference is exactly what A.11 is asked to measure.

The desensitisation is likewise a trade at an operating point rather than an architectural impossibility. Seeing past near cover and detecting a near object compete at any one setting of gain, pulse energy, range gate and exposure. A design may interleave settings and buy both, paying time and energy for the interleave. **Whether that purchase is affordable in this budget is an A.11 measurement and not something the sources settle.**

The consequence for a rule-maker is the reason this passage exists, and it is a consequence about transfer rather than about mechanism. A vendor’s product page for a Geiger-mode lidar camera states that the technology “offers the ability to confidently navigate dense foliage and other obscurants at significantly longer distances”.⁵⁷ Nothing here establishes what underlies that particular claim, and this paper does not infer it. What the published record does establish is that the demonstrated route to robust wide-area under-canopy recovery in this device class is angular diversity and multi-look fusion from a moving platform. **A capability claim read off a datasheet therefore cannot be transferred to a fixed horizontal emplacement without evidence**, and the paper records the demonstrated mechanism as one route rather than as proof of the basis of every vendor’s assertion.⁵⁸

15.5.1.2 Receiver and readout power remain substantial even when optical pulse energy is small

Single-photon sensitivity is normally presented as an energy argument: if a point can be established from a handful of photons, the transmitter can be small, and therefore—the inference runs—the

⁵⁶Marino and others (n 38), section 5, ‘Foliage penetration experiments’. The passages quoted were read from the page image.

⁵⁷BAE Systems, ‘Geiger-mode lidar camera’ product page <https://www.baesystems.com/en/product/geiger-mode-lidar-camera> accessed 27 August 2026. **Manufacturer product literature; no DOI, no independent validation, and no claim here that any figure on it has been verified.** Quoted for the vendor’s own capability statement that the camera technology “offers the ability to confidently navigate dense foliage and other obscurants at significantly longer distances”, which is the claim §15.5.1.1 examines.

⁵⁸Two further Geiger-mode foliage papers are on the open record and were **not obtained**: S E Johnson, ‘Foliage penetration optimization for Geiger-mode avalanche photodiode lidar’ in *Laser Radar Technology and Applications XVIII* (2013) Proc SPIE 8731, 87310C <https://doi.org/10.1117/12.2018911>; and J Lawrence, D B Kelley and D G Fried, ‘Ground estimation under foliage with Geiger-mode lidar’ in *Laser Radar Technology and Applications XXX* (2025) Proc SPIE <https://doi.org/10.1117/12.3053936>. **Bibliographic records verified; full texts not accessed, and no proposition in this paper rests on either.** They are recorded to show that foliage penetration is a continuing subject of open publication from 2013 to 2025, which their titles and venues establish and which is all that is claimed. No performance figure is taken from either and none should be inferred.

sensor can be small in power too. **The first half is true and the second does not follow**, because the published figures show a receive chain whose dissipation is unaffected by how few photons the transmitter sends.

On the transmit side, the foundational system’s measured transmitter characteristics are 33 μJ per pulse at 532 nm at a pulse repetition rate of 5 to 10 kHz, which is an average optical output of about **0.33 W** at the top of that range.⁵⁹ On the receive side, the published dissipation of a modern framed readout chip **alone** is **0.5 W** at 20 kframes per second with a thousand events per frame; the same source states that these architectures are built so that power and data volume are proportional to the activity on the array, by event-driven readout, region-of-interest processing, clock gating and data thinning.⁶⁰ A defence prime’s published specification for an integrated Geiger-mode lidar camera gives **18 W at saturation** for the synchronous variant and under 50 W for the asynchronous one, for a unit of 1.3 to 1.5 kg and about a litre in volume.⁶¹ Those receiver figures exclude the laser, the scanner, the pointing mechanism and the point-cloud processing.

What that comparison establishes, and what it does not. It establishes a **material receiver and readout load**, and it disposes of any inference from a small optical pulse energy to a microwatt-class complete ranging system: whatever the transmitter costs, the receive chain does not scale down with it. **It does not establish that a photon-counting architecture in general spends more energy in its receiver than in its transmitter.** The 0.33 W is average *optical output*; the receiver figures are *electrical dissipation*; and without the transmitter’s wall-plug efficiency—which none of the sources examined here reports—its electrical input is unknown and the two quantities cannot be ordered. **The paper does not order them, and an earlier version of this passage did.**

Two consequences survive that narrowing, and they are the ones the budget needs.

First, the always-on question is settled at the cited operating point. The paper’s tier-one continuous allowance is 500 μW (§17.3). The readout chip alone, at the operating point its authors give as typical, dissipates one thousand times that. **That operating point is therefore ruled out as the continuously powered tier in this budget**, which is what the architecture at §15.5.2 needs to know: the affordable arrangement is invocation on an event, for a bounded dwell, after something cheaper has decided that looking is worth the energy. It is not a proof that no photon-counting receiver, under some other architecture or operating point, could approach continuous operation—only that this one cannot, by three orders of magnitude, and that a designer proposing otherwise carries the burden of the measurement.

⁵⁹Marino and others (n 38), Table 3, ‘The laser transmitter characteristics’, giving 33 μJ per pulse at 532 nm and a pulse repetition rate of 5 to 10 kHz under a column headed *Measured Value*. The product at the top of that range, about 0.33 W, is this paper’s arithmetic. **It is an average optical output. It is not an electrical input**, and the transmitter’s wall-plug efficiency is not reported in any source examined here.

⁶⁰Aull and others (n 40): “Power dissipation of the latest framed readout chips depends on frame rate and activity level, but 0.5 W is a typical value for the R7 chip operating at 20 kframes/s with 1000 events per frame”, and the architectural statement that the designs “leverage the sparse detections to make the power and data volume proportional to the activity on the array”. **Both figures were read from the page image as well as the text layer.** The figure is for the readout chip alone and is not a system power.

⁶¹BAE Systems (n 57), specification table: synchronous camera 128 $\times$ 32 format, 50 μm pixel pitch, maximum frame rate 90 kHz, maximum range depth 6 km, range resolution 7.5 cm (500 ps), volume 900 cm^3 , weight 1.3 kg, **power at saturation 18 W**; asynchronous camera 32 $\times$ 32 format, 65.8 μm hexagonal pitch, range resolution 4.8 cm (320 ps), volume 1000 cm^3 , weight 1.5 kg, **power at saturation under 50 W. These are stated by the manufacturer for the camera; the page does not state a system power, and no system-level figure is inferred from them here.**

Second, receiver power is a function of scene activity, which couples the weather term to the energy term through a mechanism this paper had not identified. If dissipation tracks activity, and activity is what fog, rain and snow produce, then the adverse-weather result at §15.5.1 is not only a detection-performance result and not only a false-cue result. It is a *power* result: the same unwanted returns that §15.5.6 traces into dwell or into the cue stream also raise the dissipation of the receiver that registers them. P_m **is not a constant**. It is a function of the environmental state z , and it moves in the same direction as $F_{m,\text{raw}}^{\text{day}}$. A.11 measures it against activity rather than assuming a datasheet figure.

15.5.1.3 The narrow instantaneous field, and what one sector sweep costs

A ranging sensor of this class does not see a sector. It sees a small solid angle and must be pointed. The foundational system’s total instantaneous field of view is **10.3 by 10.3 milliradians**—about 0.59 by 0.59 degrees—and its published wide-angle imagery was assembled from multiple step-stare collections positioned by pan-tilt encoder readouts, covering about 8 milliradians in elevation by about 160 milliradians in azimuth.⁶² **Nine degrees of azimuth was a composite.**

The cost of covering a sector follows directly. **Solid angle, not planar area, is the right measure**, and for a centred rectangular sector of angular width θ_{az} and height θ_{el} the exact expression is

$$\Omega_{\text{rect}} = 4 \sin^{-1} \left[\sin \left(\frac{\theta_{\text{az}}}{2} \right) \sin \left(\frac{\theta_{\text{el}}}{2} \right) \right]$$

so that, with Ω_{FOV} the sensor’s total instantaneous field, an overlap factor $\kappa \geq 1$ and a per-pointing dwell t_{point} :

$$N_{\text{point}} = \kappa \frac{\Omega_S}{\Omega_{\text{FOV}}}, \quad E_{\text{sweep}} = P_m t_{\text{point}} N_{\text{point}}$$

*[Illustrative estimate, not a design point, and deliberately favourable to the sensor. A centred 60° by 10° sector is $\Omega_S = 0.17437 \text{ sr} = 572.4 \text{ deg}^2$, which is **4.6 per cent smaller than the planar product** of the two angles and is the figure used here. The published $10.3 \times 10.3 \text{ mrad}$ field is $\Omega_{\text{FOV}} = 1.0609 \times 10^{-4} \text{ sr} = 0.34827 \text{ deg}^2$. $\kappa = 1$, which no real scan achieves; $t_{\text{point}} = 10 \text{ ms}$, being a hundred frames at the published 10 kHz imaging rate. The powers are the public commercial proxy envelope of §17.3 and the published camera figure of §15.5.1.2; **they are proxy points and not demonstrated lower bounds**. The estimate excludes scan overlap, mechanical step and settle, startup, transmitter power, the pointing mechanism and point-cloud processing.]*

Quantity	Value
Pointings per sweep at $\kappa = 1$	1,644
Active time per sweep at 10 ms per pointing	16.44 s
Energy per sweep at 18 W	82.2 mWh
Energy per sweep at 29 W	132.4 mWh
Full sweeps per day that consume the entire 1.221 Wh passive budget	9.2 to 14.9
One sweep per hour at 18 W	1.97 Wh/day — 161 per cent of the passive budget

That is the quantitative form of what §15.5.2 states qualitatively, and it is why scheduled active search is classed as a different mission posture rather than as a variant: on assumptions

⁶²Marino and others (n 38), Table 1 and section 4.

chosen to favour the sensor, one sweep an hour costs more than everything else the platform does, and every omitted term—overlap, settling, startup, transmitter, processing—moves the figure the same way.

And it exposes the design variable the detector literature has actually been working on. Since $E_{\text{sweep}} \propto P_m(N_A)/N_A$ at a fixed per-pixel field, a larger array reduces the sweep cost only to the extent that receiver power grows *sub-linearly* with format. The published progression—from 4×4 to 256×256 arrays, from 350 nm to 180 nm CMOS, with event-driven readout, region-of-interest processing, clock gating and data thinning introduced expressly to reduce per-pixel dissipation—is precisely a record of making power grow sub-linearly with format.^{63 64} **Read from this paper’s side, twenty years of that work is the work of making a sector sweep affordable.** The ratio $P_m(N_A)/N_A$ is therefore the figure of merit an unattended-emplacement designer should ask a vendor for, and it is not a figure any source examined here reports.

15.5.1.4 Eye safety and the false-alarm term are coupled through the detector, which Protocol IV’s drafters had no reason to anticipate

§15.2 states the treaty position and §15.5.1 states the engineering one: Protocol IV attaches a precaution obligation to the employment of laser systems, and supplies no engineering classification. A designer who takes that obligation seriously moves towards the eye-safer regime around 1500 nm. **That choice is not free, and where its cost lands depends on the architecture.**

Silicon Geiger-mode avalanche photodiodes do not respond usefully beyond about 900 nm, so longer wavelengths require the InGaAsP/InP material system. The published comparison of figures of merit for test devices—wafer-probed, at 5 V overbias, at -20°C , about 20 μm in diameter, and reported as best results—gives a dark count rate of **10 per second for silicon against 500 per second for InGaAsP/InP designed for 1064 nm**, with reset time an order of magnitude longer and photon detection efficiency lower.⁶⁵ The camera-level literature reports that systems delivered for 1.5 μm operation obtain comparable results “with a higher dark count rate being the primary tradeoff in providing longer wavelength detection”.⁶⁶

The direction is established and the magnitude is not. A fiftyfold ratio measured on cooled wafer-probed test devices at 1064 nm is not an array-level figure at 1550 nm at ambient temperature, and this paper does not have one.

And where the extra dark counts are paid depends on where the modality sits, exactly as §15.5.6 requires for weather. Dark counts are unwanted detections and they behave like any other unwanted detection in the transfer of §15.5.6: they must survive filtering, clustering and tracking to become anything at system level, and what they become depends on the architecture.

⁶³Aull and others (n 40), on the migration from 350 nm to 180 nm CMOS and on event-driven readout with region-of-interest processing adopted to reduce per-pixel power dissipation.

⁶⁴Aull (n 41), on the same progression in array format and readout architecture.

⁶⁵Aull and others (n 40), Table I, ‘Figures of merit for test devices’, reporting for silicon a photon detection efficiency of 0.7 at 405 nm, a dark count rate of 10 per second, a reset time of 100 ns and 300 ps timing jitter; and for InGaAsP/InP designed for 1064 nm a photon detection efficiency of 0.5, a dark count rate of 500 per second, a reset time of 1 μs and 300 ps timing jitter. The table’s own conditions are stated with it and must travel with the numbers: **best results, on wafer-probed test devices of about 20 μm diameter, at 5 V overbias and -20°C .** The same paper states that silicon does not respond usefully beyond approximately 900 nm, which is why longer wavelengths require the InGaAsP/InP system.

⁶⁶Entwistle and others (n 39), on the dark-count tradeoff at 1.5 μm .

- **Where the active modality is permitted to originate or replace the tier-one cue**, surviving dark counts contribute to $F_{1,m}$ and therefore to N_3 , which §17.4 identifies as the dominant term in the energy budget. On that architecture the eye-safety decision reaches the endurance claim directly.
- **Where the modality is invoked only after a passive escalation**, they do not originate tier-one trips at all. They degrade confirmation reliability, raise the processing load, and lengthen dwell through retries—entering E_m through $t_{\text{obs},m}(z)$ rather than the cascade through F_1 .

Either way the obligation the designer discharges in Part V is paid for in Part VI, and the confirmatory architecture is again the cheaper place to pay it.

This is a coupling, not a criticism of the instrument. Protocol IV was drafted against blinding as a method of warfare and could not have been drafted against detector shot noise. The observation is offered to designers and to those conducting a legal review of a new weapon: an eye-safety decision taken late, as a compliance step after the sensing architecture is fixed, changes the dark count rate and, through whichever of the two paths above applies, the energy budget and the endurance claim. **It is a systems decision wearing the clothes of a compliance decision**, and A.11 measures it directly by running the same rig at both wavelengths where the hardware permits.

15.5.2 Four architectures, not one sensor

Most of the confusion in this area comes from treating “has lidar” as a binary property of a platform. It is not: the same device supports four architectures with different energy forms, different emission consequences and different standing in this paper.

Variant	Function	Energy form	Emission consequence	Status here
Passive EO/thermal baseline	identification after a passive cue	existing tier-three term	no active optical emission	baseline
One-time lidar survey	metric landmark and terrain model at emplacement	E_{survey}/H	brief emission during emplacement	admissible variant
Event-driven lidar	metric confirmation after a passive escalation	$\frac{N_3 \alpha_m}{3600} (E_{\text{start}} + P_m t_{\text{obs}} + E_{\text{proc}})$	emission only on selected events	principal add-on analysed
Scheduled lidar search	search beyond the passive cue footprint	$n_s E_{\text{sweep}}$, priced at §15.5.1.3	predictable or cryptographically scheduled bursts	separate mission posture
Continuous lidar search	continuous active 3D sensing	$24P_m$	persistent active emission	different architecture

The one-time survey is the easiest addition and the one this paper would make. It runs once, at the moment §15.4 already admits a ranged measurement, when a human or a carrier robot is standing at the site and the emplacement is the loudest event in the platform’s life. It improves the metric landmark set on which the survey-integrity condition depends. **It does not leave the thirty-day passive baseline untouched merely because it runs once, and an**

earlier version of this paper said that it did. At the top of the public proxy envelope a five-minute survey is $29 \text{ W} \times \frac{5}{60} \text{ h} = 2.42 \text{ Wh}$, which amortised over the thirty-day horizon is **80.6 mWh/day, about 6.6 per cent of the passive baseline**, before any allowance for pointing, startup or processing. Even a one-minute survey is about **16.1 mWh/day, or 1.3 per cent**. **If deployment equipment supplies the survey energy, the operational platform budget is unchanged, although the deployment logistics are not. If the emplacement battery supplies it, E_{survey} must be measured and amortised over H .** The system boundary must be stated either way; it is the statement, not the smallness of the number, that keeps the baseline honest.

Event-driven lidar is the easiest addition that changes performance, and is therefore the one that has to be priced and tested rather than assumed. It fires only on escalations the passive cascade has already produced, so its energy is bounded by N_3 and its emission by the escalation rate rather than by the clock.

Scheduled search is now priced and the price is the argument against it. §15.5.1.3 computes the cost of one sector sweep from a published instantaneous field of view and the public power envelope, and finds that between 9.2 and 14.9 full sweeps a day would consume the entire passive budget. The classification of scheduled search as a different mission posture is therefore not a matter of taste.

Continuous active search is a different paper. It abandons the premise on which Part IV, Part VI and the endurance claim all rest. It may well be the right design for some mission; it is not this one, and it must be costed and classified separately rather than admitted through the phrase “with lidar”.

15.5.3 The substitution problem, which is the reason this subsection exists

Consider a platform whose safeguards are implemented where safeguards are usually implemented: in the perception pipeline, as constraints on what the camera’s classifier may output and act on. A confidence floor, a class allow-list, a minimum bounding-box size, a rule that release requires a face-on human silhouette—each is a real control and each is attached to the camera.

Now substitute the modality. A lidar path can detect a candidate, measure its dimensions, classify it from the point cloud, track it across observations, and issue an effect request, without any part of that chain passing through the camera classifier or its constraints. Nothing about the platform’s name, its procurement description, its autonomy label, its effector or its model family has changed. **The controls have not been removed. They have been bypassed, because they were attached to a sensor rather than to a function.**

This is the same structural finding Part IX reaches from the historical record—that the safeguard is a layer, and a layer can be removed—arriving from the design side and one step earlier. There, the safeguard was removed by a decision. Here it is bypassed by a component substitution that nobody need characterise as a change of policy at all.

15.5.4 The authority ceiling, and the gate that enforces it

The repair is to attach the constraint to the *function* and place it downstream of every path that can reach the effector.

The authority ceiling. Let τ_H be the time of the last human decision that sets or enlarges authority, and let $A_H = A(\tau_H)$ be the resulting authority set—the ground, the target class, the

conditions and the period a person authorised. The invariant is that authority may contract but never enlarge without a further human act:

$$\forall t \geq \tau_H, \quad A(t) \subseteq A_H$$

and, for every sensing modality or fused path $m \in M$,

$$\forall m \in M, \forall t \geq \tau_H, \quad U_m(t) \subseteq A_H$$

where U_m is the usable envelope of §17.6 evaluated for that modality. **The consequence is the point of the whole subsection and should be stated as a non-implication:**

$$C_{\text{ID,lidar}} \supset C_{\text{ID,visual}} \not\Rightarrow A_{\text{lidar}} \supset A_{\text{visual}}$$

A better sensor sees further. It does not thereby acquire permission to act further. That the second does not follow from the first is obvious when written down and is exactly what a perception-attached control fails to enforce.

A stronger rule may be specified if a designer wants it—monotone contraction, $t_2 > t_1 \geq \tau_H \Rightarrow A(t_2) \subseteq A(t_1)$ until a new human act occurs—but it is **an option and not a consequence of the ceiling**, and it has a cost: it forbids automatic recovery inside the original ceiling after a temporary sensor degradation, so a platform that contracted its envelope in fog could not re-expand it when the fog cleared without a person. Which of the two to specify is a mission judgement.

The common gate. Let $y_m(x, t)$ be the perception output of modality m for an object or location x , $V(t)$ the survey-integrity predicate of §15.4, and $\Gamma(t)$ the authorisation-valid predicate of §25.3. The gate G is a single function through which every path must pass:

$$G(y_m(x, t), A_H, V(t), \Gamma(t)) = 1 \Rightarrow x \in A_H \wedge V(t) \wedge \Gamma(t)$$

equivalently, and this is the form to implement against:

$$x \notin A_H \Rightarrow G(y_m(x, t), A_H, V(t), \Gamma(t)) = 0 \quad \forall m$$

The quantifier over m is the whole content. The gate sits downstream of electro-optical, thermal, lidar and fused perception and upstream of any effect request, and it must not be implemented inside any classifier. An arrangement in which the camera pipeline enforces the constraint and the lidar pipeline reaches the effector by another route satisfies every description of the system and none of its safety argument. The architecture is the one the safety-critical control literature calls a simple, independently verifiable controller constraining a more capable one,⁶⁷ and the reason for adopting it here is that the more capable path is precisely the one whose behaviour cannot be characterised in advance.

The gate does not replace the physical safing of §26 or the independent authorisation channel of §25.3. It is a third thing, and the three are independent by construction.

⁶⁷D Seto, B H Krogh, L Sha and A Chutinan, ‘The Simplex architecture for safe online control system upgrades’ in *Proceedings of the 1998 American Control Conference* 3504–3508 <https://doi.org/10.1109/ACC.1998.703255>.

Bibliographic record verified; full text not read. Cited for the general architectural pattern of a simple, independently verifiable controller constraining a more capable one, which is what §15.5.4’s gate is; no claim is made that the cited work addresses weapons, autonomy or this application.

15.5.5 The modality materiality test

A designer, a purchaser and a reviewer all need a rule for when a sensor change is a component substitution and when it is a different system. The rule this paper proposes is a vector rather than a label:

$$\Delta_m = (\Delta P_D, \Delta P_{ID|D}, \Delta F_1, \Delta q_2, \Delta E_{\text{day}}, \Delta \Sigma, \Delta C_{\text{cue}}, \Delta C_{\text{ident}}, \Delta C_{\text{pose}})$$

A modality change is material when it changes any safety-relevant component of this vector beyond the validated bounds, even if the autonomy label, model family, platform name and effector remain unchanged. Such a change requires the relevant verification and review to be repeated.

Each component is measurable, and A.11 is the experiment that measures them. The test is offered as an engineering criterion. It has an evident affinity with the question a weapons review asks about material change in a means or method of warfare,⁶⁸ but this paper does not claim that any instrument or review guide adopts this vector, and §51 does not treat it as law.

The September 2026 texts make this a question about a word the drafters tried and then dropped. The draft report before the Group of Governmental Experts on 3 September 2026 would have had States “ensure that modifications to existing LAWS, that have a material impact on the operational functioning of their target identification, selection, and engagement functions, or on their anticipated effects, are reviewed to ensure compliance with IHL” (CCW/GGE.1/2026/CRP.1/Rev.1, paragraph 38(b)). The qualifier has a history worth recording. The rolling text of 18 December 2025 required review of further development or modifications with a “significant impact” on the same functions or effects; the chair’s summary of the first 2026 session reports that many delegations found that qualifier insufficiently clear, that alternatives such as “substantive” were proposed, and that others suggested deleting the qualifier or confining re-review to modifications that effectively result in a new type of system (CCW/GGE.1/2026/WP.2, 1 April 2026, paragraph 61); the draft of 7 July 2026 had “substantial impact”; the draft of 3 September had “material impact”; and the report the Group adopted on 4 September 2026 carries no provision on the review of modifications at all. Its paragraph 36 restates the obligation to determine the lawfulness of a new system at study, development, acquisition or adoption, and stops there (CCW/GGE.1/2026/3, paragraph 36). The vector defined above is a measurable form of the qualifier the drafters could not settle, over exactly those functions: a modality change is material if the measured change in the detection, conditional identification, false-trip, leakage, energy, emission-signature or envelope terms falls outside validated bounds, and immaterial if it does not. Whether a treaty text should carry a measurable test, or any re-review obligation, is a question for drafters. The engineering observation is only that “material impact” admits of one, that A.11 is the experiment that would calibrate it, and that the adopted text now reaches a sensor change only through the review of a new system—the gap the interval companion identifies at its Part IV.

15.5.6 Where the modality enters, which decides what the weather costs

The proposition that snow “manufactures tier-one false trips” is **conditional, and the condition is architectural**. It is true only where the active sensor is permitted to originate or replace the tier-one cue. Where lidar is invoked only after a passive escalation, weather appears in the budget through dwell, retries, candidate load and confirmation performance—not as a new stream of wakes.

⁶⁸International Committee of the Red Cross, *A Guide to the Legal Review of New Weapons, Means and Methods of Warfare* (ICRC 2006). **Not read in full.** Cited only for the general proposition that a review regime attaches significance to material change in a means or method of warfare. **No claim is made that the Guide adopts, anticipates or endorses the m vector**, which is this paper’s engineering proposal.

Where lidar sits	What weather actually does
One-time deployment survey	Survey corruption, reduced usable range, pose error, or a delayed emplacement. No recurring tier-one wake stream , because the sensor is not running during the emplacement.
Event-driven tier-three confirmation	More candidates, longer dwell, retries and possible identification errors inside an already-invoked active window . No additional tier-one wake arises from those returns.
Scheduled active search	Weather returns can create a modality-specific cue stream during scheduled on-time .
Continuous search, or lidar replacing the passive cue	Weather returns can directly contribute a new tier-one false-cue stream , materially altering both energy and false-escalation rates.

The transfer, written down. Let z be the weather state and $\lambda_m(x, z)$ the density of unwanted raw detections at location x , in detections per second per square metre. Over an active on-time $\tau_{\text{on},m}$ per day and an observed support A_m :

$$F_{m,\text{raw}}^{\text{day}}(z) = \tau_{\text{on},m} \int_{A_m} \lambda_m(x, z) dA$$

That quantity is not a count of system alerts and must not be used as one. An unwanted raw detection becomes a system-level cue only if it survives sensor filtering, clustering, temporal-consistency checks and tracking, so

$$F_{1,m}(z, \theta) = q_{\text{survive},m}(z, \theta) F_{m,\text{raw}}^{\text{day}}(z)$$

where $q_{\text{survive},m}$ is that survival probability. **It is unknown until measured, and A.11 measures it.** The supplement’s F1_modality refuses to compute without it rather than substituting a plausible value.

Where the active modality is permitted to cue, the cascade separates by modality:

$$N_3 = Tp_1p_2 + F_{1,\text{passive}}q_{2,\text{passive}} + F_{1,m}q_{2,m}$$

Where it is event-driven only, $F_{1,m}$ does not enter the tier-one cascade at all, and the weather burden appears instead in the observation duration carried back into E_m :

$$t_{\text{obs},m}(z) = t_{m,0} + n_{\text{retry}}(z)t_{\text{retry}}$$

The design consequence is worth stating because it is the reason to prefer the confirmatory architecture. An event-driven active tier inherits the weather burden as *dwell*, which is bounded by the escalation rate the passive cascade already controls. A cueing active tier inherits it as *rate*, which is **not bounded by the passive cascade**—scheduled search remains bounded by its scheduled on-time and continuous search by the clock, but neither is bounded by N_3 , which is the quantity the rest of this budget is built on. The paper’s recommendation at §15.5.2 that lidar be admitted as a one-time survey or as an event-driven confirmation, and not as a search sensor, therefore has a second justification independent of emission: **it is the arrangement in which bad weather cannot open a new term in the budget.**

15.6 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 §15.4 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.

15.7 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.

And a common gate downstream of every perception path, which §15.5.4 specifies and which is a third thing rather than a restatement of either. Taken together the release condition is compact:

$$\text{Release}(x, t) \Rightarrow V(t) \wedge \Gamma(t) \wedge x \in A_H \wedge x \in C_{\text{effect}}$$

—survey integrity holds, authorisation is live, the object lies inside the authority a person conferred, and the effector can reach it. **The implication runs one way only.** Satisfying all four does not require release; it is the minimum without which release must not occur, and the physical safing above remains independent of every term in it.

Why the second channel is independent of the sensing should be stated, because fielded practice now does the opposite. In a machine-vision weapon station under contract for a counter-drone role, the reporting describes detection, tracking and the firing solution as computed by the station

from a passive camera suite, in autonomous or semi-autonomous modes, with a person required to approve each engagement.⁶⁹ That arrangement is workable while the person is present, and as a safeguard it has a property this specification declines to inherit: the function that could hold fire to spare a human being lives inside the same perception function that finds the target—one substrate, one code path, one set of parameters—so the safeguard is coupled to the most modifiable subsystem of the machine and shares its failure modes. A safety property implemented inside the system it is meant to constrain is a single point of failure in the ordinary engineering sense, and a rule that the property must be present does not change that. The two interlocks above are the alternative: physical safing that no software state can alter, and an authorisation channel that shares nothing with the sensing and inference chain, gated by the displacement check at §15.4 and bounded by the expiry at §26. A safeguard that is architecturally independent of the perception it restrains is the only kind that is not a configuration choice. Removing it produces a different machine rather than the same machine with a setting changed, which is what “not trivially removable” means as a specification. §34 records the commercial systems whose vendors state that the trigger is the only thing a person holds.

The clearest instance in the public record is not a robot at all but a static armed turret, and it is worth setting out because it is the manufacturer’s own account rather than an inference from reporting. Of a sentry turret which its maker states detects and tracks to about three kilometres, a senior engineer of the company said in 2015: “It wasn’t initially designed this way. Our original version had an auto-firing system. But all of our customers asked for safeguards to be implemented.” He added that this “wasn’t a problem for us” technologically, and that the customers “were concerned the gun might make a mistake.”⁷⁰ The safeguard there is, on the maker’s account, a configuration retrofitted to an auto-firing baseline at a customer’s request, and its presence was described as a commercial choice rather than a technical constraint. Part IX §49 assembles the wider record of which that is one case, and the pattern it finds is the reason this specification does not place release inside the perception chain: across the fielded systems whose sequence is documented, the autonomous capability came first and the human control was added afterwards, as a mode.

15.8 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

⁶⁹Defence Blog, ‘U.S. Marine Corps taps an AI-powered gun against cheap drones’ (21 July 2026) <https://defence-blog.com/u-s-marine-corps-taps-a-ai-powered-gun-against-cheap-drones/> accessed 24 August 2026, describing the Bullfrog weapon station as relying on a fully passive camera sensor suite for detection and tracking, operating in autonomous or semi-autonomous modes, and still requiring a person to approve any engagement, under a prototype agreement of approximately USD 6.2 million. The earlier feasibility-study award for the same station is at §32.3.

⁷⁰Simon Parkin, ‘Killer robots: The soldiers that never sleep’ (*BBC Future*, 16 July 2015), quoting a senior research engineer of DoDAAM on the Super aEgis II; reproduced with attribution in Erico Guizzo and Evan Ackerman, ‘Do We Want Robot Warriors to Decide Who Lives or Dies?’ (*IEEE Spectrum*, 31 May 2016) <https://spectrum.ieee.org/do-we-want-robot-warriors-to-decide-who-lives-or-dies> accessed 26 August 2026, which also records the manufacturer’s statement that its systems as configured will not fire until a person confirms the target. The quotations are taken from the IEEE reproduction, which credits the original; see Part IX §49.

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.

15.9 Communication

Default posture 1, receive only, per Part IV §13. 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 §11. 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.

15.10 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.⁷¹ 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.⁷² 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.

15.11 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. §15.6 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 §15 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.

⁷¹Rivero (n 5).

⁷²*ibid*, giving approximately USD 36,000 for mobile repair vehicles developed by the Lithuanian Riflemen's Union.

Part VI—Energy

16. Endurance metrics and the holding mission

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.

16.1 What the quoted figure measures

Quoted endurance is a drive-hours figure. Fielded platforms are rated at four to eight hours for small types and eight to sixteen hours for mid-sized logistics types.⁷³ The stated ambition for the next generation is endurance measured in days rather than hours.⁷⁴

A second, independent data point comes from the other side of the technology. 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.⁷⁵ Two states, two morphologies, two industrial bases, and the same order of magnitude. The reconnaissance figure 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 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 quoted figure is therefore not an upper bound on emplaced duration. §17.2 gives the comparison in fielded terms: a thirty-day unattended sensor system against an eight-hour portable surveillance radar, both published, both fielded.

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, and its Defence Minister announced in April a target of 25,000 unmanned ground vehicles to be 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.⁷⁶ The platform codified under NATO

⁷³Unmanned Systems UK (n 2).

⁷⁴*ibid.*

⁷⁵Army Recognition (n 21). These are figures reported from a state demonstration and are treated on the same footing as the Ukrainian reporting at (n 4).

⁷⁶Katie 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 16 October 2026, reporting the target, announced by the Defence Minister after a meeting with domestic manufacturers, that 25,000 unmanned ground vehicles be contracted in the first half of 2026; his statement of 18 April 2026 that more than nine thousand mis-

cataloguing standards in the same month is specified as a 300-kilogram payload with a 50-kilometre range.⁷⁷

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 §16.3 has to propose one rather than adopt one.

The first reported case of ground robots taking a position without infantry, in Kharkiv Oblast in the summer of 2025, was an assault rather than a hold, and it is described in exactly those terms.⁷⁸ 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.

16.2 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 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 §17.3 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.

sions were carried out using unmanned ground vehicles in March alone; and the stated objective that all frontline logistics be performed by robotic systems.

⁷⁷ibid, reporting the codification of the Bizon-L, a 300-kilogram-payload logistics platform with a 50-kilometre range, under NATO cataloguing standards.

⁷⁸ibid, reporting the President’s address of 14 April 2026, which described how, in the summer of 2025, 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, and called it the first capture of an enemy position by unmanned platforms alone in the war. The operation is described throughout as the capture of a position rather than its retention.

16.3 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. §17.3 states it for the design point—fifty seconds—and gives the budget as a function of it, so that the cost of a stricter posture can be read off rather than guessed.

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.

16.4 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.

17. The architectural inversion

17.1 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.⁷⁹ 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.⁸⁰

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.**

17.2 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**.⁸¹ Set against that a man-portable ground surveillance radar of the same era: a national forces manual gives the internal-battery endurance of the AN/PPS-15(V)2 as **sixteen hours** at 21 °C.⁸²

Set those two side by side. Same task—detecting movement across ground. Same era of electronics. **Thirty days against sixteen hours is a factor of about forty-five, 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. The two figures are drawn from different documents about different systems and are not a controlled comparison; what they support is the order of magnitude and its cause, and the derivation at §17.3 is what carries the argument.

⁷⁹The 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.

⁸⁰Multi-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.

⁸¹On 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/sys/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.

⁸²Department of National Defence (Canada), *Ground Surveillance Radar*, B-GL-317-015/FP-001 (CFP317(15)), giving the internal battery endurance of the AN/PPS-15(V)2 as “16 hours at 21°C (70°F)” on a BA 4386/PRC-25 battery https://radionerds.com/images/e/e8/PPS-15_B-GL-317-015-FP-001.pdf accessed 26 August 2026. In v1.0.0 and v2.0.0 this note carried a comparison with the AN/PPS-5D quoted at approximately eight hours; that figure could not be resolved to any primary source and has been withdrawn, and the sourced figure for a man-portable radar of the same class is given in its place. Appendix A.7 records the correction.

That is this paper’s central claim, evidenced on fielded hardware by the organisations that fielded it, decades ago. The inversion at §17.1 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. This paper’s budget assumes two hundred tier-one trips a day, yielding ten tier-three escalations. The fielded system is specified to survive thirty days while *detecting* a thousand events a day—five times the trip rate assumed here—which suggests the assumption is conservative rather than optimistic at the cheap tier. §17.3 varies that rate to two thousand a day and the thirty-day requirement remains about 290 Wh, so the calibration is reassuring rather than threatening.

And a current commercial system brackets this paper’s own budget from the other side. A manufacturer’s published description of a fielded seismic unattended ground sensor family states a **one-month mission life** for its tactical node, **powered by three AA lithium cells**, and a mission life of **more than two years** for its long-term node; it gives detection of a person’s movement to about 100 m and of light vehicles to about 300 m; and it describes an imaging node whose “[c]amera [is] activated only upon sensor detection, conserving battery life”.⁸³ A second paper from the same manufacturer gives the core system’s power as approximately **10 mW** and its transmit power as **15 mW**, and describes the system as “operationally secure, incorporating features that provide low-probability of intercept and low-probability of detection”.⁸⁴

Three of those figures bear directly on Part VI and none of them was available to this paper before. The **thirty-day emplacement on a handful of primary cells** is the design case at §15 fielded as a product, which makes the horizon at §17.3 a commercial specification rather than an extrapolation. The **imager woken only on a sensor detection** is the tiered architecture of §17.1 in a catalogue, which means the inversion this Part proposes for an armed platform is already the ordinary arrangement for an unarmed one. And the **core at approximately 10 mW** is the first published comparator for the always-on floor this paper assumes at 500 μ W plus 2 mW of housekeeping: the fielded figure is roughly four times higher, which is what one would expect of a node that also self-locates and maintains a mesh, and it is the right order. These are a manufacturer’s own published figures and are treated as such; they are not an independent measurement, and §22 continues to say that nothing in this Part has been measured on a bench by the author.

⁸³Textron Systems, ‘MicroObserver UGS System’ product literature, supplied to this compilation as a portable document and read in full. The document states one-month mission life for the MO-1045 node powered by three AA lithium batteries, greater than two-year mission life for the MO-2730 node, detection of a person’s movement to a distance of up to 100 metres and of light vehicles to up to 300 metres, gateway-to-sensor standoff of one to five kilometres with options to ten, a programmable hibernate mode, and an imager node whose camera is “activated only upon sensor detection, conserving battery life”. These are a manufacturer’s published figures for its own product and are not an independent measurement.

⁸⁴William Agassounon and Robert Steadman (Textron Defense Systems), *The Application of Unattended Ground Sensors to Human Socio-Cultural Behavior Monitoring* (17 October 2012), DTIC AD-A581145, giving approximately 10 mW for the core system, 15 mW transmit power, low-probability-of-intercept and low-probability-of-detection features, and describing the product as having a “negligible false alarm rate, and extremely long battery life (2+ years)”. A manufacturer’s account of its own product.

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 kind §17.1 describes.⁸⁵ 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 §17.3. The forty-five-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.

17.3 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 §15.1. 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. What has changed since the first version of this Part is how it is presented: every line below is *derived* from a single stated parameter vector rather than assumed independently, which is what makes the budget checkable and what exposed an inconsistency in the earlier presentation. Appendix A.8 records the correction and the reproduction procedure.

The parameter vector. Eleven quantities fix the cascade, which is the part of the budget that moves. They are stated once, here. **They are not the whole budget:** the receiver, housekeeping and effect lines are separately stated assumptions, set out beneath the table and again at A.8, and the honest form of the claim is that *eleven cascade parameters, together with separately stated receiver, housekeeping and effect assumptions, determine the budget*. Every one of the seventeen is either a component class from a published figure or an assumption, and each is labelled where it is used; none is a measurement made by this paper.

Symbol	Quantity	Design point	Class
P_1	tier-one continuous power	500 μ W	design allowance; acoustic subassembly anchored, integration unmeasured
T	genuine events entering the sector, per day	2	assumed
F_1	tier-one false trips per day	200	assumed
p_1	probability tier one registers a genuine event	0.95	assumed
p_2	probability tier two confirms a genuine event	0.95	assumed
q_2	probability tier two passes a tier-one false trip	0.041	assumed

⁸⁵On 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.

Symbol	Quantity	Design point	Class
P_2	tier-two power when live	50 mW	component class
t_2	tier-two live time per wake	10 s	assumed
P_3	tier-three power when live	10 W	component class ⁸⁶
t_3	tier-three live time per escalation	30 s	assumed
t_b	tier-three settle and resume time, at full power	10 s	assumed

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

The two event rates follow from the vector rather than being chosen. Tier two is woken by every tier-one trip, genuine or false; tier three is woken only by what tier two passes:

$$N_2 = Tp_1 + F_1 = 201.9 \text{ wakes per day}$$

$$N_3 = Tp_1p_2 + F_1q_2 = 1.81 + 8.20 = 10.0 \text{ escalations per day}$$

The decomposition of N_3 is the result §17.4 is about and it should be read off here: **82 per cent of the escalations are false-alarm leakage and 18 per cent are genuine.**

The lines. The daily total is the sum of three fixed loads and three driven by the event rates:

$$E_{\text{day}} = 24P_1 + N_2P_2t_2 + N_3P_3(t_3 + t_b) + E_{\text{rx}} + 24P_{\text{hk}} + \frac{n_e P_e t_e}{H} + E_m$$

with powers in watts and times in hours, E_{rx} the receiver line, P_{hk} housekeeping, and the effect term n_e engagements of $P_e t_e$ each amortised over a horizon H of thirty days. **E_m is the active-modality term of §15.5.2. It is zero in the passive baseline and otherwise takes one of four non-zero forms, which are not interchangeable:**

Architecture	E_m	Daily total form	Status
Passive EO/thermal baseline	0	1.221 Wh/day	costed here

⁸⁶On 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.

Architecture	E_m	Daily total form	Status
One-time survey	E_{survey}/H	1.221 Wh/day + E_{survey}/H	priced at \$15.5.2, not assumed negligible; excluded from the platform budget only on a stated system boundary
Event-driven active	$\frac{N_3 \alpha_m}{3600} (E_{\text{start},m} + P_m t_{\text{obs},m} + E_{\text{proc},m})$	scales with the escalation rate	priced below against a public proxy envelope; E_{start} , E_{proc} , t_{obs} and α_m are measured at A.11
Scheduled active search	$n_s E_{\text{sweep}}$, with E_{sweep} priced at \$15.5.1.3	dominated by the sweep rate	different mission posture; one sweep an hour exceeds the whole passive budget
Continuous active search	$24P_m$	dominated by P_m	different architecture; outside this budget

[Symbolic in P_m , $t_{\text{obs},m}$ and α_m , which remain variables outside the explicitly labelled sensitivity below. **No military-grade active-sensor design point is asserted**, because none is public. What can be said without measurement is structural: a one-time survey is amortised over the horizon but is **not therefore negligible**—at the top of the proxy envelope five minutes is 6.6 per cent of the daily budget and one minute is 1.3 per cent (\$15.5.2)—an event-driven branch is bounded by N_3 and therefore by the false-alarm rate that already dominates the budget, and scheduled and continuous branches are **not bounded by the passive cascade at all**: scheduled search is bounded by its scheduled on-time and continuous search by the twenty-four hours in a day, which at 20 W is 480 Wh, about four hundred times the whole passive budget. E_{survey} is in joules where the expression divides by 3600, and in watt-hours where it does not; the code states which at each use.]

A public commercial proxy envelope, to expose the scale of the omitted term without manufacturing a specification. No military active-sensor design point is asserted. The table below evaluates a **public commercial proxy envelope of 6.5–29 W**, taken from three manufacturers’ own published electrical-power figures for commercial sensors.^{87 88 89} These are manufacturer specifications, not independent measurements, not peer-reviewed figures, not military-system estimates, and not claims about range in this paper’s geometry. Observation duration, invocation fraction, startup energy and processing energy remain architecture-specific variables to be measured under A.11. **The resulting calculations are sensitivities, not predictions of fielded-system consumption.**

[*Illustrative sensitivity, not a design point.* $N_3 = 10$ events/day, $\alpha_m = 1$, one-second active observation, and **startup and incremental processing energy excluded**. Percentages are of the 1.2213 Wh/day passive baseline. Derived; the proxy powers are manufacturer-reported.]

⁸⁷Livox (n 46), for the lower point of the public commercial proxy envelope.

⁸⁸Ouster (n 47), for the middle band of the public commercial proxy envelope.

⁸⁹Hesai Technology, ‘OT128 product specifications’ <https://www.hesaitech.com/product/ot128/> accessed 27 August 2026, stating power consumption of 29 W. **Manufacturer specification; not an independent measurement, not peer-reviewed, and not a military design point.**

Public proxy power	One-second contribution	Per cent of the 1.221 Wh/day baseline
6.5 W	18.1 mWh/day	1.48%
10 W	27.8 mWh/day	2.27%
20 W	55.6 mWh/day	4.55%
29 W	80.6 mWh/day	6.60%

A ten-second observation multiplies each of those by ten, giving approximately **181 to 806 mWh/day** across the same envelope—between a seventh and two-thirds of the entire passive budget—and those values **still exclude startup and incremental processing energy**.

Which is why the invocation expression separates three terms rather than two. Warm-up, stabilisation, scan time, point-cloud processing, retries and shutdown are not one quantity, and folding them into $P_m t_{\text{obs}}$ hides two of the three:

$$E_m = \frac{N_3 \alpha_m}{3600} (E_{\text{start},m} + P_m t_{\text{obs},m} + E_{\text{proc},m})$$

with $E_{\text{start},m}$ and $E_{\text{proc},m}$ in joules per invocation, P_m in watts, $t_{\text{obs},m}$ in seconds, and the divisor converting joules per day to watt-hours per day. For the sensitivity form above, which temporarily omits the two unmeasured terms, $E_m[\text{mWh/day}] = 2.7778 \alpha_m P_m[\text{W}] t_{\text{obs},m}[\text{s}]$ at $N_3 = 10$.

The two omitted terms are not small, and that is the point of separating them. *The following two values are illustrative counterfactuals chosen solely to show the algebraic sensitivity.* They are not component figures, not vendor data, and not estimates: no source examined here reports a startup or processing energy for any device of this class. At 20 W for one second, a startup and a processing cost of 20 J each would take the daily figure from 55.6 to 167 mWh; at 100 J each, to 611 mWh—more than eleven times the observation term alone. **Neither term has been measured**, by this paper or in any source it has found, and A.11 measures both. A budget that reports only $P_m t_{\text{obs}}$ has reported the term least likely to dominate.

Evaluated at the vector above, with $E_m = 0$:

Line	Derivation	Per emplaced-day
Tier one, continuous	500 $\mu\text{W} \times 24 \text{ h}$	12.0 mWh
Tier two, woken	50 mW $\times 10 \text{ s} \times 201.9$	28.0 mWh
Tier three, woken	10 W $\times (30 + 10) \text{ s} \times 10.0$	1,111.7 mWh
Communication, scheduled listening	30 mW $\times 1 \text{ s}$, once every 60 s	12.0 mWh
Communication, message reception	30 mW $\times 200 \text{ s} \times 3 \text{ messages}$	5.0 mWh
Housekeeping	2 mW $\times 24 \text{ h}$	48.0 mWh
Effect	20 W $\times 5 \text{ s} \times 5 \text{ engagements} \div 30 \text{ days}$	4.6 mWh

One tier-three escalation costs 111 mWh. The receive chain is two lines rather than one because the recall link, computed at Appendix A.9, closes only on a receiver held up during a scheduled window: the standing cost is the listening, and the event-driven cost is the reception of an authenticated command, which §25.2 shows takes two hundred seconds rather than one.

Total: 1.221 Wh per emplaced-day, of which tier three is **91.0 per cent**. Over thirty days, **36.6 Wh**.

A published anchor for the tier-one order of magnitude, and what it is not. The 500 μW is the largest unsourced input in this budget and A.2 exists to close it. It is not, however, an unsupported invention. A 2023 review of acoustic wake-up technology reports **a 305.5 μW wake-up chip—300 μW for the MEMS microphone and 5.5 μW for the signal classification circuit—for acoustic recognition of tracked and wheeled vehicles.**⁹⁰ The review’s own framing is the qualification that matters: it groups that device with “low-power acoustic wake-up microchips *without back-end function modules*”, which is to say a chip and not a complete wake-up microsystem. The figure therefore anchors the electronic order of magnitude of an always-live acoustic front end, and it does **not** validate an integrated seismic-and-acoustic tier, detection of a walking person at 30 m, the probability of detection this paper’s cascade assumes, or the assumed false-trip rate. The underlying device paper has not been obtained and no figure is taken from it beyond what the review states. A separately published VLSI implementation of an energy-aware acoustic wake-up detector for a surveillance sensor network establishes that the architecture has been built and reported in the peer-reviewed record; its full text has not been read and no figure is taken from it.⁹¹

The tier-one accounting boundary, stated so the anchor cannot be misread as the whole tier. What §17.3 charges as P_1 is a sum:

$$P_1 = P_{\text{WU, audio}} + P_{\text{WU, seismic}} + P_{\text{gate/leak}}$$

Against the 500 μW design allowance, and taking the published figure at face value for the acoustic term alone:

$$P_{\text{WU, audio}} = 305.5 \mu\text{W}, \quad P_{\text{WU, seismic}} + P_{\text{gate/leak}} = 194.5 \mu\text{W}$$

The first term is a published subassembly figure reported in a review; the second is an unmeasured integration allowance and not a measurement of anything. It is stated this way so that a reader can see exactly how much of the tier is anchored and how much is stipulated, and so that the 305.5 μW device is not mistaken for a functional equivalent of the tier this paper specifies.

One boundary must be drawn before the bench measurement, or it will be double-counted. A.2 calls for an integrated front end measured *with* its microcontroller, real-time clock

⁹⁰Deng Yang and Jiahao Zhao, ‘Acoustic wake-up technology for microsystems: a review’ (2023) 14(1) *Micromachines* 129 <https://doi.org/10.3390/mi14010129>. **Full text examined** via the publisher’s open-access PDF. The review states, in section 4.1: “A 305.5 μW wake-up chip, 300 μW for the MEMS microphone and 5.5 μW for the signal classification circuit, is reported for acoustic recognition of the tracked vehicle and wheeled vehicle”, and introduces that passage by noting that “some low-power acoustic wake-up microchips **without back-end function modules** are reported, too”. **This is a secondary report:** the review is the source relied on here, the underlying device paper has not been obtained, and no figure is taken beyond what the review states. It anchors an order of magnitude for an always-live acoustic front end and **validates neither an integrated seismic-and-acoustic tier, nor pedestrian detection at 30 m, nor any probability of detection or false-trip rate assumed in this paper.**

⁹¹David H Goldberg, Andreas G Andreou, Pedro Julián, Philippe O Pouliquen, Laurence Riddle and Rich Rosasco, ‘VLSI implementation of an energy-aware wake-up detector for an acoustic surveillance sensor network’ (2006) 2(4) *ACM Transactions on Sensor Networks* 594–611 <https://doi.org/10.1145/1218556.1218562>. **Bibliographic record verified against Crossref; full text not accessed**, the publisher’s platform returning an error to automated retrieval, which was not circumvented. Cited only for the proposition its title and record establish—that an energy-aware acoustic wake-up detector for a surveillance sensor network has been implemented in VLSI and reported in the peer-reviewed literature. **No power figure, detection range or field result is taken from it and none should be inferred.**

and conditioning, while the budget above charges a separate 2 mW housekeeping line. The supervisor and the real-time clock belong in one place or the other. **Either include them in P_1 and exclude them from P_{hk} , or measure them within P_{hk} and exclude them from P_1 ; and where the bench yields only a combined idle-rail figure, charge that figure once.** A.2 is amended accordingly.

And the conclusion does not turn on which value is right, which is worth demonstrating rather than asserting.

[Derived from the model at A.10 by varying P_1 alone, every other input held at the design point. The first row takes the published acoustic subassembly figure as though it were the whole tier, which it is not; the third doubles the design allowance as a stress case.]

P_1 case	Tier-one energy	Total daily budget	Thirty days
305.5 μ W — acoustic subassembly only	7.332 mWh/day	1.2167 Wh/day	36.50 Wh
500 μW — the design allowance	12.000 mWh/day	1.2213 Wh/day	36.64 Wh
1 mW — stress case	24.000 mWh/day	1.2333 Wh/day	37.00 Wh

Doubling the assumed tier-one load changes the total budget by less than one per cent. The exact value of P_1 is therefore not load-bearing for the endurance conclusion, and a reader who thinks 500 μ W optimistic may substitute a milliwatt and reach the same thirty-day result. **A.2 is still necessary, and for a different reason than power.** What is missing is not a wattage but the coupled surface—detection range against probability of detection against false-trip rate against power—and it is the false-trip rate, not the tier-one draw, that §17.4 shows to dominate.

The readiness level, stated with the figure as §16.3 requires. The interval from a tier-one trip to a release decision being available is $t_2 + t_b + t_3 = \mathbf{50 \text{ seconds}}$. The metric at §16.3 is meaningless without it, and it is not free: readiness and energy trade against one another directly, and a commander choosing a posture is buying response time by the watt-hour.

Readiness interval	Total	Discrimination line	Multiple
20 s	0.51 Wh/day	417 mWh	5.4×
30 s	0.79 Wh/day	695 mWh	8.3×
50 s	1.22 Wh/day	1,112 mWh	11.1×
70 s	1.78 Wh/day	1,668 mWh	16.2×
95 s	2.35 Wh/day	2,223 mWh	19.0×

Four things about that number deserve attention, and the third is a requirement rather than an observation.

Effect is not the driver. The weapon costs 4.6 mWh a day against roughly 1,152 mWh 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 assembled 3-to-20 W spread: *[Derived.]*




Figure 1. The readiness–energy trade. Each point is a row of the readiness table at §17.3: the daily total at a stated tier-two live time, gate settle time and tier-three live time, plotted against their sum—the interval from a tier-one trip to a release decision being available. The relation is close to linear across the range a commander would consider, which is the useful form of the result: response time is bought by the watt-hour at an approximately constant rate.

Tier-three power	Total per emplaced-day	Over 30 days	Tier-three share
3 W	0.44 Wh	13.3 Wh	75.3%
10 W	1.22 Wh	36.6 Wh	91.0%
20 W	2.33 Wh	70.0 Wh	95.3%

At the pessimistic end the thirty-day requirement is 70.0 Wh, about 280 g of cells, and the December harvesting requirement rises to 6.3 Wp — roughly 0.038 m². Neither is a difficulty. The architecture absorbs a factor of nearly seven in the largest line of its own budget.

Tier three must be gated to zero, and this is a specification requirement rather than a convenience. The budget above assumes the tier draws its 10 W only while live and *nothing* between escalations. That assumption is doing more work than any other in the Part, because an inference module held in a retention state instead of being fully unpowered costs its retention current continuously: *[Derived.]*

Tier-three retention power	Total	As a multiple of the gated design
0 mW (gated)	1.22 Wh/day	1.00×
10 mW	1.46 Wh/day	1.20×
100 mW	3.62 Wh/day	2.97×
300 mW	8.42 Wh/day	6.90×

A hundred milliwatts of retention triples the budget and three hundred multiplies it by seven. **The architecture therefore requires either wake-on-event silicon that genuinely idles at zero, or a hard power gate with a microcontroller-class wake stage in front of it; a general-purpose module suspended rather than gated defeats the design.** The t_b term above is what a hard gate costs on each escalation, and it is included precisely so that the

alternative can be priced against it. This is the assumption on which a hostile embedded-systems reader should press first, and it is stated here so that the pressing is easy.

Tier three dominates, so the event rate is the sensitive parameter. Varying the false-trip rate with q_2 held: *[Derived.]*

F_1 per day	N_3 per day	Total	Over 30 days
50	3.9	0.52 Wh	15.5 Wh
200	10.0	1.22 Wh	36.6 Wh
800	34.6	4.04 Wh	121.1 Wh
2,000	83.8	9.67 Wh	290.1 Wh

Even at two thousand false trips a day—ten times the design point and twice the rate the fielded comparison at §17.2 was specified to survive—the thirty-day requirement is 290 Wh, about 1.2 kg of cells. 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 $15 \text{ W} \times 24 \text{ h} = \mathbf{360 \text{ Wh per day}}$, and **10.8 kWh over thirty days**. The ratio at the design point is **295:1**. At the 800-trip rate it is 89:1. 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.

17.4 Where the event rate comes from, and how much of the result depends on it

§17.3 derives N_3 from the vector rather than assuming it. This subsection states what that buys, because the escalation rate governs the whole budget and a reader is entitled to know how much of the conclusion rides on quantities the paper has assumed rather than measured.

Let the sector receive T genuine events of interest per day. Let tier one produce F_1 false trips per day from wind, precipitation, animals, distant traffic and sensor noise. Let p_1 and p_2 be the probabilities that tiers one and two register a genuine event, and let q_2 be the probability that tier two passes a tier-one false trip rather than rejecting it. Then

$$N_3 = Tp_1p_2 + F_1q_2$$

and the daily energy is the fixed loads plus N_3 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.]*

At the design point the two terms are 1.81 and 8.20. **Eighty-two per cent of the escalations that dominate the budget are false-alarm leakage**, and that is the result this subsection exists to produce.

Which parameters actually matter. Differentiating the total with respect to each parameter in turn and normalising gives the local elasticity of the budget to that parameter—a first-order sensitivity at the design point, not a variance decomposition over the ranges, and the third column

below normalises those elasticities rather than apportioning variance—the percentage change in the total produced by a one per cent change in the parameter:

$$\varepsilon_x = \frac{\partial E_{\text{day}}}{\partial x} \cdot \frac{x}{E_{\text{day}}}$$

[Derived.]

Parameter	Local elasticity	Running share of the summed elasticities
P_3 , tier-three power	+0.91	23%
F_1 , false-trip rate	+0.77	43%
q_2 , tier-two leakage	+0.75	62%
t_3 , tier-three live time	+0.69	79%
t_b , settle and resume	+0.23	85%
T , p_1 , p_2	+0.16 each	97%
housekeeping	+0.04	98%
P_2 , t_2 , receiver power, P_1	+0.02, +0.02, +0.01, +0.01	100%

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

The environmental term has been measured, at scale, and the measurement is old. The largest unattended-sensor deployment ever mounted supplies it. A study of that campaign records that the sensors were first tested in Panama, with the intent of replicating the conditions expected in the theatre of use, and that in the theatre “[t]he presence of frogs, animals and thunder caused a sensor activation rate **three to four times higher** than in Panama”, producing “491,814 ADSID activations and 125,649 acoustic acoubuoy activations” during the initial deployment.⁹² A NATO technical report on a later generation of the same class of system records the same problem in its own words: “While high detection rates resulted, the systems were plagued by very high false alarm rates”, among the causes being “[w]ildlife (which causes false alarms on fielded sensors)”.⁹³

Those are the two propositions §17.3’s F_1 term asserts, evidenced. A change of environment alone moved the activation rate by a factor of three to four, on the same hardware and against the same intended targets—which is the quantified form of the claim below that a sector’s cost is a property of its weather and its ground. The absolute counts are not comparable to this paper’s design point, being a theatre-scale deployment over years rather than one emplacement over thirty days; they are given because they establish the order of the problem rather than its rate.

And the remedy that campaign adopted is the one §17.3’s cascade is built on. The same study records that sensors were combined “to limit the number of false detections”, and that “[b]y requiring the sensor to detect the same target in multiple ways before transmitting a ‘hit,’ the chances of a false report were reduced”—one combination requiring “both a Magnetic and seismic

⁹²Eric D Haider, *Unattended Ground Sensors and Precision Engagement*, Master’s thesis, Naval Postgraduate School, Monterey (December 1998), DTIC AD-A359912, at the section headed ‘False Activations’. The activation counts are quoted by that thesis from its own cited source. The deployment concerned is the sensor campaign in Southeast Asia; the figures are theatre-scale and over an extended period, and are given here for the order of the problem and for the ratio between environments, not as a rate comparable to this paper’s design case.

⁹³NATO Science and Technology Organization, *Disposable Multi-Sensor Unattended Ground Sensor Systems for Detecting Personnel*, STO Technical Report TR-SET-158, AC/323(SET-158)TP/578, DTIC AD-A626493, read in full. The report gives no numeric false-alarm rate in its text and the quoted statements are qualitative; they are cited as an institutional characterisation of the problem and not as a measurement.

detection before reporting”.⁹⁴ That is tier-one-gates-tier-two, fielded under operational pressure fifty years before this paper proposed it for an armed platform, and adopted for exactly the reason the elasticity table gives: because rejecting false trips in the cheap tiers is the highest-leverage thing available.

The energy budget of a holding system is set by its false-alarm rate, not by enemy activity. F_1 and q_2 enter as a product and carry equal elasticity, together nearly two-fifths of the whole sensitivity, while the genuine-event terms T , p_1 and p_2 carry a fifth between them. Wherever F_1 greatly exceeds T —the normal condition for an unattended sensor, since the environment is busy and the enemy is not—the second term dominates. 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 available at low cost. Tier-three power has the largest elasticity, but halving it means a better imager and processor. Halving q_2 has almost the same effect and means a better magnetometer and passive-infrared channel, which are cheap. That is where the engineering effort belongs.

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 §15.1 raises F_1 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.

How much of the conclusion survives the uncertainty. Every figure above is a point estimate, and a point estimate is not an adequate answer to the question a technical adviser will ask. Sampling all ten varying parameters simultaneously across the ranges this Part states—tier-three power 3 to 20 W, live time 20 to 45 s, settle 5 to 20 s, false-trip rate 100 to 400 a day, q_2 from 0.020 to 0.082, genuine events 1 to 5 a day, and the fixed loads across their component classes—with a triangular distribution on each and its mode at the design point, over 200,000 draws: *[Derived; the procedure is at Appendix A.8 and the distribution is plotted at Figure 2.]*

Quantity	5th percentile	Median	95th percentile
Total, Wh per emplaced-day	0.80	1.75	3.59
Identification and release judgement, mWh/day	684	1,628	3,454
Discrimination multiple	6.6×	14.2×	28.7×

Two readings of that table matter more than the interval itself.

The architecture survives throughout the modelled support. At the ninety-fifth percentile the thirty-day requirement is 108 Wh—about 430 g of cells, or a harvesting surface of roughly 0.058 m², still smaller than a sheet of A4. Nowhere in the sampled ranges does the energy budget

⁹⁴Haider (n 92), at the section headed ‘Combined Sensors’.




Figure 2. The sampled distribution of the daily total. The design point sits below the median because several parameters are bounded further above the design value than below it, which is a property of the stated ranges and not a hedge. No part of the sampled space places the thirty-day requirement above about 110 Wh. That is a statement about the modelled support under independent triangular sampling, not about every design this specification admits.

defeat the design. That is a statement about the ranges this Part states and the distribution it assumes over them, not a proof about the whole parameter space: the sampling is independent across parameters, the triangular shape is a judgement rather than a measurement, and a reader who believes F_1 and q_2 are positively correlated—which is plausible, since a noisy site degrades both—will obtain a wider upper tail than this. That is a materially stronger statement than the point estimate, and it is the one a sceptical reader should be given.

And the discrimination multiple never falls below about six and a half. The claim that identification and release judgement cost roughly an order of magnitude more than holding ground without them is robust across the sampled range: it is $6.6\times$ at the fifth percentile and $29\times$ at the ninety-fifth. The *number* at \$21.2 is a point estimate and should be quoted as one; the *proposition* that discrimination is the dominant cost, and dominant by about an order of magnitude, is not sensitive to any assumption in this Part.

One calibration of the vector should be recorded against the current literature, and it does not flatter the paper. A 2023 survey of seismic target detection reports classification accuracies for personnel-against-vehicle discrimination in the range of about 88 to 93 per cent across the studies it reviews, and states the false-alarm mechanism directly: ambient vibration may suppress the signature sought, and “excitations from non-footstep sources may resemble footprint excitations, increasing the likelihood of false alarms”.⁹⁵ This paper assumes $p_1 = p_2 = 0.95$. That is at or above the optimistic end of what the survey reports, and a reader substituting 0.90 will find the genuine-event term falls by about a tenth while the budget barely moves, because the budget is

⁹⁵Priyankar Choudhary, Neeraj Goel and Mukesh Saini (Indian Institute of Technology Ropar), ‘A Survey on Seismic Sensor based Target Detection, Localization, Identification, and Activity Recognition’ (2023) *ACM Computing Surveys*, doi 10.1145/3568671, read in full. The accuracies quoted are those the survey reports for individual studies it reviews and are not a single measured figure.

dominated by the false-alarm term and not by this one. The assumption is optimistic and it is not load-bearing, which is the honest way to state it.

Five modality-dependent sensitivities belong on the same list and none of them can yet be filled in. If an active tier is added, the parameters that move are α_m (the fraction of escalations that invoke it), P_m and t_m (its power and dwell), F_1 and q_2 (which it may raise or lower), and the environmental vector z , through which all of the preceding become weather-dependent in a way the passive budget is not. **The temptation to assume that a metric sensor lowers the false-alarm rate should be resisted: it is a hypothesis, and the measured evidence at §15.5.1 runs partly the other way**, since snowflake backscatter in a coaxial active configuration is reported to generate false alarms at a rate no passive seismic front end approaches.⁹⁶ Whether an active confirmation tier reduces q_2 enough to pay for E_m is exactly what A.11 measures, and it is a different answer in every z .

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, F₁ and q₂ 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.

17.5 The energy-constrained detection range, and the usable field envelope

The previous subsection records that detection range and energy cost are coupled and pull against each other. That was left as an observation. It can be made into a quantity, and doing so produces the most practically useful result in this Part, because it answers a question every specification for this class of system gets wrong: *how far can the thing see?*

The question is malformed, and the reason is the cascade. Raising tier-one sensitivity extends the range at which a genuine event is registered and, in the same motion, raises the rate at which the environment trips the detector. The false trips are what the budget spends on—§17.4’s elasticities put F_1 and q_2 together at nearly two-fifths of the whole sensitivity—so extending nominal range spends energy at a rate set by the environment rather than by the enemy. **A front end therefore has two ranges: the one at which it can detect, and the one at which it can be afforded to detect for thirty days.** How far apart they are is an empirical question about a particular sensor, site and threshold family; that they are different quantities is not.

Definition. Let $m \in M$ be the sensing modality or fused sensor-path state, $\theta_m \in \Theta_m$ the thresholds and configuration for that modality, z the environmental condition vector, and r range. Write $P_D(r \mid m, \theta_m, z)$ for the probability of detection, $P_{ID|D}(r \mid m, \theta_m, z)$ for the conditional probability of the required identification given detection, $F_1(m, \theta_m, z)$ for the false-trip rate, $q_2(m, \theta_m, z)$ for leakage into the expensive tier, E^* for the daily energy ceiling and S^* for the permitted emission-signature set. The daily budget is the cascade of §17.3 with the environment-dependent terms made explicit:

$$E_{\text{day}}(m, \theta_m, z) = 24P_1 + N_2P_2t_2 + N_3P_3(t_3 + t_b) + E_{\text{rx}} + 24P_{\text{hk}} + E_{\text{effect}} + E_m$$

⁹⁶Rasshofer, Spies and Spies (n 45), on snowflake backscatter generating a false-alarm rate in a coaxial active configuration.

with $N_2 = Tp_1 + F_1$ and $N_3 = Tp_1p_2 + F_1q_2$ as before, all durations in hours, and E_m the modality term defined at §15.5.2. The **energy-constrained detection range** is then

$$R_E(m, z; E^*, S^*) = \sup_{\theta_m \in \Theta_m} \sup \{ r \geq 0 : P_D(r \mid m, \theta_m, z) \geq p_D^*, P_{ID|D}(r \mid m, \theta_m, z) \geq p_{ID}^*, \\ F_1(m, \theta_m, z) \leq F^*, E_{\text{day}}(m, \theta_m, z) \leq E^*, \Sigma_m \in S^* \}$$

The outer supremum over θ_m is not decoration. An earlier version left the threshold free inside the set, which makes the expression ill-formed: the range achievable at *some* threshold is not the range achievable at the threshold the designer must actually choose. The quantity is the best range obtainable over the admissible configurations, subject to every constraint holding at that same configuration.

Two further points of construction. **Detection and identification are separate constraints and must not be interchanged:** where identification is not required at the cue stage, define a cue-only $R_{E,D}$ by dropping the second condition, and where the joint event matters write $P_{D \cap ID} = P_D P_{ID|D}$. And **emission is a constraint, not a cost:** $\Sigma_m \in S^*$ is a separate admissibility test precisely because collapsing emission into energy is the error Part II exists to prevent.

What the definition is, and what it is not. It is a constrained specification quantity together with the experiment that measures its inputs. **A numerical value of R_E is not a result until P_D , $P_{ID|D}$, the false-alarm rate, the energy and the environment have been measured together, on the same hardware, at the same threshold.** Appendix A.2 is that experiment for the passive baseline and A.11 for an active modality. Everything below the definition is a worked example on stipulated relations.

A hypothetical worked example, on stipulated relations. [*ILLUSTRATIVE MODEL — NOT FIELD DATA.*] To show the shape of the constraint the two sensing relations must be parameterised, and neither is measured. Signal level is taken to fall with the transmission loss the cited seismic study gives—a geometric term for a Rayleigh surface wave plus an excess term of about 0.30 dB per metre at 100 Hz—with a 6 dB spread for gait, coupling and site; the false-trip rate is taken to fall by a decade for each 10 dB the threshold is raised. **Both are stipulations.** The source level is calibrated so that at the design threshold the range at which the assumed detection law reaches $P_D = 0.95$ is the 30 m the field trials at §15.1 report as reliable footstep detection—which is a *calibration to a literature figure, not a demonstration that the literature figure corresponds to $P_D = 0.95$* , since the cited trials do not state a detection probability. [*Derived from the cascade at §17.3; the two sensing relations are illustrative and stated above.*]

Threshold	Range at $P_D = 0.95$	F_1 per day	N_3 per day	Total	Over 30 days
60 dB	2 m	2	1.9	0.29 Wh	8.8 Wh
50 dB	11 m	20	2.6	0.38 Wh	11.3 Wh
40 dB	30 m	200	10.0	1.22 Wh	36.6 Wh
30 dB	55 m	2,000	83.8	9.67 Wh	290 Wh
20 dB	82 m	20,000	822	94.2 Wh	2,825 Wh
10 dB	111 m	200,000	8,202	939 Wh	28,175 Wh

Three things follow **within the illustrative model**, and each is a statement about the model rather than about any real front end.




Figure 3. ILLUSTRATIVE MODEL — NOT FIELD DATA. The energy-constrained detection range under the stipulated sensing and false-alarm relations of this subsection. Each point on the curve is a tier-one threshold: moving right lowers the threshold, extends the range and admits more of the environment. The horizontal line is the budget at \$17.3. The two ranges the system possesses are the intersection with that line and the vertical asymptote at the front end's own limit.

The two ranges the model produces differ by a factor of about four. At the design budget the worked example detects to 30 m; the same assumed front end at its own limit reaches 111 m, admitting two hundred thousand false trips a day and costing 28.2 kWh over thirty days—two and a half times what the conventionally architected platform of \$17.3 costs in the same period. In the model, the energy argument does not merely fail at that range; it inverts.

Range is bought at a steeply diminishing rate. In the model, doubling the energy budget moves the range by about eight metres and doubling it again by about eight more; across a factor of eight in energy, and in cells, and in harvesting area, the range moves from 20 m to 46 m.

The threshold is an endurance decision. Whoever chooses θ_m is choosing the emplacement's lifetime, and in the model a 10 dB error in that choice moves the thirty-day requirement by a factor of eight.

Under the illustrative sensing and false-alarm relations used here, the curve has these values. The contribution is the definition and the measurement procedure, not the numerical range. Whether a real front end exhibits a fourfold gap, or a twofold one, or none, is an empirical question about a particular sensor, site, passband and threshold family, and Appendix A.2 is the experiment that would determine it. What is not model-dependent is the *structure*: sensitivity and false-alarm rate move together, false trips drive the expensive tier, and therefore a sensing range and an affordable sensing range are different quantities that a specification must state separately. **No specification this paper has examined states the threshold, and none states the false-trip rate it implies.** Both belong next to any range figure for an emplaced sensing weapon, and that recommendation does not depend on the illustrative curve at all.

17.6 The usable field envelope, and what “closed” means at formation scale

The same argument generalises, and stating it generally is what makes the metric at §16.3 usable by someone specifying a real system rather than this one.

A fielded system does not possess the range of its longest-range component. It possesses the region in which it can simultaneously cue, identify, know that its surveyed geometry still holds, still be authorised to act, and physically reach—while inside its energy and emission budgets. Per modality, and at a time t because authority moves:

$$U_m(t) = C_{\text{cue},m} \cap C_{\text{ident},m} \cap C_{\text{pose}} \cap A(t) \cap C_{\text{effect}}$$

The intersection is the primary form because real sectors are not circular: terrain, cover, the surveyed landmark set and the arc a weapon is bolted to are all shaped, and the region where all five hold is shaped by all of them. A radial simplification is available and is **only** a simplification:

$$R_{\text{usable},m} = \min(R_{\text{cue},m}, R_{\text{ID},m}, R_{\text{pose}}, R_{\text{authority}}, R_{\text{effect}})$$

Each term is a constraint this paper has established, and naming them together is the contribution:

Term	What bounds it	Where	Modality-dependent
$C_{\text{cue},m}$	the energy-constrained detection range $R_E(m, z)$	§17.5	yes
$C_{\text{ident},m}$	the range at which the tier can classify well enough to release	§15.1, §21.2	yes
C_{pose}	where the survey-integrity condition $V(t)$ holds	§15.4	partly
$A(t)$	the authority ceiling: surveyed sector, horizon, live authorisation	§15.6, §25.3	no, by construction
C_{effect}	what the weapon can physically reach	§15.1	no

The fourth row is the one §15.5 exists to protect. Every other term moves when the sensor changes. Authority does not, because it is conferred by a person and not measured by an instrument.

17.6.1 Four envelopes, and where closure actually sits

An earlier version of this paper said that any externally supplied cue makes the architecture open. **That was wrong, and it contradicted §6.1, §7 and §15.1.** The withdrawal test asks whether a party *other than the operator* can cause the mission to fail by withdrawing a service. A cue arriving over the operator’s own private mesh from the operator’s own second emplacement fails no part of that test: nobody outside the operator’s control can withdraw it. What the cue does cost is **platform independence**, which is a different property and worth naming separately rather than collapsing into closure.

Envelope	Meaning	Closure status
Intrinsic effect envelope	what the effector can physically reach	not a sensing-system property
Platform-supported envelope	what one emplacement can cue, identify and authorise by itself	platform-independent and closed
Operator-group-supported envelope	what the operator's own private emplaced group can support, with no third-party service	closed , but not platform-independent
Third-party-enabled envelope	the additional region requiring a supplier, proprietor or outside service	open on the withdrawal test

The third row is where a real formation operates and where §15.1's distributed sensing already sits. The fourth is what Part III audits. **The distinction that matters legally and operationally is the third-to-fourth boundary, not the second-to-third**, and an earlier version drew it in the wrong place.

17.6.2 Two independence fractions, and two coverage measures

Writing $\mu(\cdot)$ for a measure on the region—area for a sector, or authorised frontage where the mission is stated as frontage—the share of the weapon's reach that is genuinely supported divides in two:

$$\eta_{\text{platform}} = \frac{\mu(U_{\text{platform}})}{\mu(C_{\text{effect}})}, \quad \eta_{\text{closed group}} = \frac{\mu(U_{\text{closed group}})}{\mu(C_{\text{effect}})}$$

Both are dimensionless and $\eta_{\text{platform}} \leq \eta_{\text{closed group}} \leq 1$. **A system with $\eta_{\text{closed group}}$ near one is what it appears to be; a system with $\eta_{\text{closed group}}$ near zero is a networked weapon with a sensor attached, and the sensor is not what makes it work.** Reporting both separates a design that needs its own group from a design that needs somebody else's service, and only the second is open.

Where the paper claims a formation-level result, the natural normalisation is coverage per unit energy:

$$\Phi_A = \frac{\mu(U)}{E_{\text{day}}} \quad [\text{m}^2 \text{ Wh}^{-1} \text{ day}], \quad \Phi_F = \frac{L(U)}{E_{\text{day}}} \quad [\text{m Wh}^{-1} \text{ day}]$$

where $L(U)$ is authorised frontage. **The units must be stated every time these are reported**, because a coverage-per-watt-hour figure without them is unfalsifiable, and because Φ_A and Φ_F answer different questions: area matters for denial of ground, frontage for denial of an approach.

The consequence for a ranged effector. *Effector reach beyond the usable envelope has no independently supported mission value.* A weapon that reaches two kilometres on a platform that cues itself to thirty metres is not a two-kilometre system; it is a thirty-metre system carrying a two-kilometre weapon. The extra reach is real and can be used—on a cue from the operator's own group, which remains closed, or on a third party's cue, which does not—but the three cases must be priced and labelled separately, and a specification that reports only the effector's reach has reported the least informative of the four numbers in the table above.

17.7 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 36.6 Wh is about 147 g. *[Derived from the figures above.]* Even at the eight-hundred-trip rate, 121 Wh is under 500 g, and at the ninety-fifth percentile of §17.4 the 108 Wh is about 430 g. 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.

18. The interlock

18.1 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.⁹⁷ Winter yield at 50°N is a small fraction of the annual figure. *[Derived, and the assumption is stated: taking December at 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 1.22 Wh daily load it requires approximately **3.3 Wp**, which at typical module power density is of the order of 0.020 m²—a surface smaller than a sheet of A5 paper. *[Derived from the 370 mWh/Wp/day figure above.]* At the ninety-fifth percentile of §17.4, 3.59 Wh per day, it requires about 9.7 Wp, or roughly 0.058 m², which is still smaller than a sheet of A4. 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.

⁹⁷PVKnowHow, ‘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.

18.2 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.⁹⁸ 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 89 mWh per day, or **2.7 Wh across thirty days**. At 710 Wh/kg that is 3.8 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.⁹⁹ That is not a complication 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 89 mWh per day, or 2.7 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.

⁹⁸Tadiran 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.

⁹⁹On 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-principles-passivation-and-iot-hybrid-power-solution/> accessed 5 August 2026.

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.

19. Energy delivery

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.¹⁰⁰ 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 §16.4.¹⁰¹

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.¹⁰²

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 §18—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.

¹⁰⁰Phihong, ‘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.

¹⁰¹Delta-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.

¹⁰²The 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 Myrnohrad, carrying loads to about 200 kg per mission, across nearly 24,500 missions in the first quarter of 2026 alone.¹⁰³ 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.

20. 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.

20.1 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 capacity in 2026 is stated at more than eight million first-person-view drones a year.¹⁰⁴ Russia is reported to be planning production of more than seven million units in 2026.¹⁰⁵ Ukraine's planned procurement for 2025 alone was 4.5 million units.¹⁰⁶

¹⁰³Rivero (n 5). 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.

¹⁰⁴National Security and Defence Council of Ukraine, 'Results of Ukraine's Defense Industry in 2025: FPV Drones' <https://www.rnbo.gov.ua/en/Diialnist/7370.html> accessed 5 August 2026.

¹⁰⁵'Russia 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.

¹⁰⁶'Ukraine 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.

[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 §17.3, 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 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.

20.2 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.

20.3 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.

21. 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.

21.1 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. It performs no sensing and no judgement.

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.¹⁰⁷

One qualification belongs here before the pricing, and it sharpens the argument rather than softening it. Sensorless is not the same as indiscriminating. Emplaced munitions carrying signature classifiers have been fielded since the 1970s: an encapsulated torpedo mine in service from 1979 is described as having discriminated hostile submarines from surface vessels and from friendly submarines, and a wide-area munition delivered from 1999 used acoustic and seismic sensing to detect, track and classify, applying a mass threshold and ranking candidates before selecting one.¹⁰⁸ Part IX §44 sets out that record and distinguishes it from the influence fuzing—a magnetic threshold, a trip wire—which is what the scatterable mine families of the same era actually used. So discrimination *between classes of materiel, by signature* was available, and at emplaced-munition prices.

What was never available in such a device, at any price, is discrimination in the sense the law of targeting requires: distinguishing a combatant from a civilian, which is a judgement about a person and not a signature match against a vehicle. That is the obligation the mine law does not impose, that is what §21.2 prices at about 1.1 Wh per emplaced-day, and the distinction between the two senses is the reason the price is worth computing.

That body of law has already answered, for the sensorless 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 remotely delivered **anti-personnel** 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; the same specification attaches to non-remotely-delivered anti-personnel mines used outside a marked and monitored perimeter. Remotely delivered mines that are *not* anti-personnel are required only to carry self-destruction or self-neutralization “to the extent feasible”, with no stated period and no stated reliability. **The hard numbers exist for the anti-personnel case and nowhere else**, which matters for §51: an anti-materiel system of the class one state’s own directive exempts would fall on the side of the instrument that has no figures on it.¹⁰⁹ It further requires mines produced after 1 January 1997

¹⁰⁷Convention 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.

¹⁰⁸See Part IX §44 and the sources there, in particular Federation of American Scientists, ‘M93 HORNET’ <https://man.fas.org/dod-101/sys/land/docs/m93.htm> accessed 26 August 2026, and the descriptions of the Mk 60 CAPTOR cited at that section.

¹⁰⁹Amended Protocol II (n 107), 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.

to contain at least eight grammes of iron or equivalent, so that they can be found by ordinary detection equipment.¹¹⁰

The thirty-day figure is where this paper’s design case comes from. The thirty-day emplacement at Part V §15 is a stipulation adopted from the instrument’s Technical Annex, not a finding derived from any operational report, and it is adopted because the law has already fixed it as the period after which an emplaced effect may not be left without a mechanism to end it. The law reached it in 1996 for a device with no sensing at all.

Set the two side by side and the relationship is plain. **The mine performs no sensing and exercises no judgement. The architecture at Part V spends 1.22 Wh per emplaced-day and exercises some. The energy is what the judgement costs.**

21.2 The exchange rate, stated

The budget at §17.3 can be re-read as a price list, and re-reading it that way is more informative than reading it as a constraint. The figures below use the corrected receive chain of §25.1, which prices a link that has been shown to close; the heartbeat variant of §25.3 is given where it differs. *[Derived from §17.3 by reallocation; no new inputs.]*

Capability purchased	Cost per emplaced-day	Share
Presence detection only—knowing something is there	12.0 mWh	1.0%
Coarse discrimination—personnel against vehicle, direction, count	28.0 mWh	2.3%
Identification-and-release-processing line	1,112 mWh	91.0%
Recallability—the receive channel	17.0 mWh (6.7 mWh on the heartbeat of §25.3)	1.4% (0.6%)
Housekeeping	48.0 mWh	3.9%
Effect—the weapon itself	4.6 mWh	0.4%
Total	1,221 mWh	

Every figure is quoted at the fifty-second readiness interval of §17.3, without which §16.3 says the quantity is meaningless, and each is a point estimate whose ninetyeth-percentile interval is at §17.4.

The 1.112 Wh line is the energy allocated to identification and release processing, not a measurement of discrimination quality. A system can spend that energy and still classify badly. The energy figure must therefore be reported together with P_D , $P_{ID|D}$, the relevant false-positive or false-release rate, the tested environment and the decision threshold. **Nothing in this table is evidence that any legal obligation of distinction is satisfied**, and the reading below that discrimination dominates the budget is a statement about where the watt-hours go, not about whether what they buy is good enough.

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

¹¹⁰ibid, 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.

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 **110 mWh** per emplaced-day. A system that holds ground *and looks first* costs **about eleven times** that. Across the sampled uncertainty of §17.4 the multiple ranges from about six and a half to about twenty-nine and never approaches unity. The entire energy problem in this class of system is the cost of looking before acting.

Recallability is one part in seventy, and it was one part in a thousand until the link budget was computed. This is the finding at Part VII §25 arriving from the opposite direction, and the two agree because both follow from a receiver being cheap and a transmitter being expensive. The correction at §25.1—the always-listening stage cannot hear the control station at the design standoff, so recall runs on a scheduled sensitive receiver instead—raised the figure by an order of magnitude and left the conclusion where it was. 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.

The weapon is 0.4 per cent. Restating §17.3 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 sensorless limit of the same design. Remove the sensing and the judgement 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 obligations Amended Protocol II already imposes on the sensorless case can be priced in the energy-spending one, and the instrument distinguishes three mechanisms which the pricing has to keep apart. A self-destruction mechanism is “an incorporated or externally attached automatically-functioning mechanism which secures the destruction of the munition into which it is incorporated or to which it is attached”. A self-neutralization mechanism is “an incorporated automatically-functioning mechanism which renders inoperable the munition into which it is incorporated”. Self-deactivating means “automatically rendering a munition inoperable by means of the irreversible exhaustion of a component, for example, a battery, that is essential to the operation of the munition”.¹¹¹

What the architecture at Part V supplies is self-neutralization: an incorporated mechanism that renders the platform inoperable. It costs approximately nothing. Self-destruction is also available, and it is also cheap. Self-deactivation is the one obligation this architecture works against. The paper’s entire energy argument is directed at preventing a component from exhausting: it sizes a primary cell to outlast the emplacement several times over, it harvests, it exchanges packs, and it scavenges. A system engineered so that its battery never runs down cannot self-deactivate in the instrument’s sense. A harvested or exchanged system must therefore supply that function through a separate, irreversibly-exhausting component rather than through its main energy path. 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,

¹¹¹Amended Protocol II (n 107), arts 2(10), 2(11) and 2(12). None of the three definitions is expressed in terms of a timer or a set period; the time periods are in paragraph 3 of the Technical Annex.

because it cannot be met by a device with no sensing—costs about 1.1 Wh 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 a little over 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.

21.3 Why this reframing matters beyond this paper

Four 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 about 1.1 Wh per emplaced-day, at a fifty-second readiness interval. 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.

The fourth concerns the **direction of the proliferation risk**, and it is the one a debate organised around state programmes points the wrong way. The table prices discrimination at 91 per cent of the budget, and the ordering, if not the proportion, holds in engineering effort and in price: the discriminating tier is where the training data, the inference and the scarce skill go, and the remaining tenth is what a platform, a passive sensor and an actuator cost—items that §34 records as on sale to the public. A device that does not discriminate is therefore not the more advanced artefact but the simpler one. It is what remains when the expensive tier is absent, and the safeguard that makes the system acceptable is the hardest thing in it to build. The obligation to discriminate has to be imposed rather than assumed because the market supplies its absence more cheaply than its presence.

22. Remaining problems

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 §18.2 and, above that baseline, the exchange route at §19, 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 §17.3 and §18.1 is derived from the parameter vector at §17.3 and from published component classes. None is a bench measurement, and §17.4 states the interval within which the result sits given the stated ranges. 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—The Costs of Closure

23. The risks

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 §15.3 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 §15.8 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 §15.4 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. §15.4’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. Assessment of mitigations

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 §15.6. 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—§15.7. 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—§15.8. 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—§15.1. 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—§19. 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. The receive-only channel

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 no intended emission to distinguish it from a system with no link at all, while remaining commandable; §12 states the one qualification, and the tuned-radio-frequency wake-up stage at §25.1 is the architecture that meets it.

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 what expiry into inertness at §26 achieves over the whole emplacement and what the heartbeat at §25.3 achieves within it, and it is why these 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 §15.8 insists on the key ceremony.

25.1 What the recall link costs, and the correction the link budget forces

Earlier versions of this Part tiered the receiver as the sensing chain is tiered: an always-listening wake-up stage in the microwatt class, gating a capable receiver that is otherwise unpowered. Wake-up radio is a mature technique with its own survey literature,¹¹² dedicated wake-up receivers operate in the microwatt class, and average consumption of about 3 μ W has been demonstrated at a wake-

¹¹²Rajeev Piyare and others, ‘Ultra Low Power Wake-Up Radios: A Hardware and Networking Survey’ (2017) 19 *IEEE Communications Surveys & Tutorials* 2117.

up packet detection latency of tens of milliseconds.¹¹³ On those figures the receive chain cost about 0.9 mWh a day, and this paper said so.

The link budget at Appendix A.9 refutes that arrangement at the stated range, and the refutation is this Part’s most consequential correction. Reported wake-up receivers achieve sensitivity of the order of -90 dBm.¹¹⁴ Computed over a two-ray plane-earth model rather than a free-space one, a -90 dBm receiver on a concealed antenna half a metre above ground is reached at **1.9 km by a twenty-watt station and 4.2 km by five hundred watts**, which is not a covert station. The design case at §15 puts the control station five to fifteen kilometres to the rear. **The always-listening stage cannot hear the station at that range at any power the mission permits, so it cannot be what makes the platform recallable.** The 0.9 mWh figure was the price of a channel that does not close, and it is withdrawn.

What closes is the arrangement §15.9 already specified for a different reason: a sensitive receiver, unpowered between windows, brought up on a schedule that forms part of the cryptographic material. At -120 dBm the same twenty watts reaches **10.6 km** and a hundred watts **15.8 km**. The scheduled window was presented as a power-saving measure. It is the mechanism by which recall works at all beyond about two kilometres.

The corrected receive chain, therefore, has the same two-line shape as every other tier in this design—a standing cost and an event-driven one. *[Derived; the receiver’s 30 mW and the message length are stated below and at A.9.]*

Line	Derivation	Per emplaced-day
Scheduled listening	$30 \text{ mW} \times 1 \text{ s}$, once every 60 s	12.0 mWh
Message reception	$30 \text{ mW} \times 200 \text{ s} \times 3$ messages	5.0 mWh
Receive chain		17.0 mWh

That is **1.4 per cent of the emplaced-day budget**, against the one part in a thousand this paper previously claimed. The claim was wrong by an order of magnitude and the correction is recorded at A.7. **The conclusion it was offered in support of survives intact and should be stated in its corrected form: recall costs about one part in seventy of the budget, a third of what housekeeping costs, and one and a half per cent of what discrimination costs.** A designer who omits recall is still not economising on anything. What has changed is that the figure is now the price of a channel that has been shown to close.

The wake-up stage is retained, and it is worth keeping for what it can actually do. Within about two kilometres it is reached by a hand-held or vehicle-borne transmitter, which is the geometry of emplacement, of local safing by the force that owns the ground, and of recovery. **The platform therefore has two receive paths with different reaches and different owners: a close-range path that anyone forward can use, and a scheduled long-range path that only the station can use.** Neither costs enough to argue about, and the tuned-radio-frequency wake-up design remains the right choice for the close path because, having no local oscillator, it presents none of the unintended emission §12 describes.

¹¹³ibid; 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 μW at a detection latency of tens of milliseconds.

¹¹⁴A ‘Tuned-RF Duty-Cycled Wake-Up Receiver with -90 dBm Sensitivity’ (2018) 18 *Sensors* 86.

25.2 The message, and the law that couples timeliness to jam resistance

The recall channel's real specification is not its power. It is the *message*, and the message settles both the latency and the jam resistance, which turn out to be the same quantity seen from two directions.

What must be in a recall command. A command that an adversary can forge is worse than no command, and a command an adversary can record and replay is worse still, because replay is how a system is held armed by someone who does not own it. A minimally defensible authenticated command therefore carries a platform address, a command word, a monotonic counter so that a recorded message cannot be replayed, and an authentication tag long enough that forging one is not a search: *[Derived; the field widths are conventional minima and a designer may substitute their own.]*

Field	Bits
Platform address	32
Command word	8
Monotonic counter, anti-replay	32
Authentication tag	128
Total	200

*[Illustrative. **The 200 bits are payload.** A real message must also bind the authorisation to a platform or group, an authority version, and a validity interval, and must carry acquisition preamble, synchronisation and coding on top of all of it. Each addition lengthens the message and, at a fixed rate, lengthens the delivery time; the coupling below is unaffected because it depends on the length, whatever the length turns out to be.]*




Figure 4. What recall timeliness costs in jam resistance. Each point on a curve is a choice of information rate: reading right buys processing gain and pays for it in latency, reading left buys latency and pays for it in the distance from which a jammer can deny the channel. The link is the five-kilometre case of Appendix A.9 with a twenty-watt station and a 10 dB required margin.

The coupling. Processing gain against a jammer is $G_p = 10 \log_{10}(B/R)$ for a spread bandwidth B and an information rate R ; the time to deliver an L -bit message at that rate is $t_{\text{msg}} = L/R$. Eliminating R between them gives the result this subsection exists to state:

$$G_p = 10 \log_{10} \left(\frac{B}{L} \right) + 10 \log_{10} t_{\text{msg}}$$

For a fixed bandwidth and a fixed message, **jam resistance and delivery time are one quantity: every doubling of the time a recall message takes to arrive buys 3.01 dB of resistance to the jammer trying to stop it.** Two qualifications belong with that statement and neither disturbs it. The relation is exact only in the idealised processing-gain model, which credits the receiver with the full bandwidth-to-rate ratio and no implementation loss; a real receiver gives up several decibels to synchronisation, imperfect despreading and analogue-to-digital headroom, and a designer should carry an implementation margin. And the two hundred bits are **payload**: acquisition preamble, synchronisation and any forward error correction sit on top of them, so the delivery times below are lower bounds on a real packet and the processing gains are upper bounds on a real link. The *slope*—3.01 dB per doubling—is unaffected by either, because both add a constant. At $B = 1$ MHz and $L = 200$ bits the constant term is 37.0 dB, and the schedule follows: *[Derived.]*

Information rate	Delivery time	Processing gain
200 bit/s	1 s	37 dB
20 bit/s	10 s	47 dB
5 bit/s	40 s	53 dB
1 bit/s	200 s	60 dB
0.2 bit/s	1,000 s	67 dB

Two things follow and both were wrong in earlier versions of this paper.

Sixty decibels of processing gain costs three and a half minutes, not one second. A one-second detection window is long enough to *find* a spread carrier at 60 dB of gain, and this paper previously conflated finding the carrier with receiving the message. It is not long enough to deliver two hundred authenticated bits at the rate that gain requires. The receiver must therefore stay up once it detects a signal, which is what the five milliwatt-hours in the table above pays for.

Worst-case recall latency is the window interval plus the message, and at the design point that is about four minutes. Sixty seconds to reach the next window, two hundred to receive the command: **260 seconds.** Not sixty. A designer who wants recall in ten seconds must raise the rate to twenty bits per second, and pay for it with 13 dB of processing gain—which, on the jamming arithmetic at A.9.3, multiplies the standoff from which a jammer can deny the channel by about a factor of two. **The recall channel cannot be made both fast and hard to jam, and the exchange rate is 3.01 dB per doubling.** That is the arithmetic behind the phrase “in a timely manner”, and §51 takes it to the text.

25.3 The heartbeat inversion, which is the architecture this analysis recommends

Everything above prices a channel that carries a command to *stop*. That is the arrangement the treaty text assumes, it is the arrangement this paper assumed, and on the paper’s own findings it has a defect that no amount of receiver design repairs.

The defect is stated at §27.1 and it is uncomfortable: **on a command-to-stop channel** a receive-only system is denied by jamming without being located, jamming removes the operator’s ability to stop it and removes nothing else, and therefore jamming makes such a system more dangerous rather than less. §25’s own list of limits says the design must be structured so that the absence of a message is the safe state. A stop-command channel does the opposite: absence of a message *is* the armed state.

The repair is to invert the channel. Require a periodic authenticated permission to continue, rather than listen for a command to stop. The platform arms only while an authenticated *continue* message keeps arriving on the schedule, and safes itself when one is missed. Nothing else in the specification changes: the same receiver, the same waveform, the same key material, the same windows. What changes is the sense of the message and therefore the direction of every failure.

The two architectures differ in exactly one respect and in every consequence, and the table is the compact statement of this Part. *[Derived; the energy figures are from §25.1 and the table below.]*

Condition	Command-to-stop	Permission-to-continue
Ordinary operation	armed; a stop command may arrive	armed; authorisation keeps arriving
Benign link outage	stays armed indefinitely	safes within one heartbeat period
Deliberate jamming	stays armed, and cannot be stopped	safes within one heartbeat period
Receiver failure on the platform	stays armed, undetectably	safes within one heartbeat period
Operator killed, withdrawn or overrun	stays armed until its own expiry	safes within one heartbeat period
What the operator knows	nothing; no confirmation is possible	nothing about <i>this</i> moment, but a bound on how long the platform can still be armed
What the adversary gains by denying the channel	a weapon that cannot be recalled	a weapon that stops
Receive chain	17.0 mWh/day	6.7 mWh/day at a six-hourly period

The right-hand column is not better in every row—rows two through five are an operational cost, not a gain—but **the failures all move to the same side of the ledger, and it is the side where the cost is borne by the force that emplaced the system rather than by whoever walks into it.**

Five consequences follow, and the first is the one the rest of this paper has been unable to obtain by any other means.

Jamming safes the system instead of protecting it. The adversary’s cheapest attack on the channel becomes the operator’s safety mechanism, and the perverse incentive §27.1 identifies for the command-to-stop case—that an adversary who understands the architecture has an active reason *not* to jam it—is removed, because jamming and clearing now point the same way. This is the single strongest argument for the inversion, and it is obtained **at no additional platform cost relative to maintaining the receive channel at all**: the heartbeat replaces the command channel’s standing listening cost rather than adding to it, and at a six-hourly period it is cheaper than what it replaces.

Partial, stale or unauthenticated messages must not reset the timer. The predicate $\Gamma(t)$ is true only on a message that authenticates, whose counter advances, and whose validity interval covers the present moment; a fragment that fails any of those leaves the timer running. This is the clause an implementer will be tempted to relax under field conditions, because a platform that safes itself on a marginal link is an operational nuisance, and relaxing it converts the architecture back into the one it replaced. $\Gamma(t)$ **is also one of the four inputs to the common release gate at §15.5.4**, so a relaxation here propagates into every perception path at once.

The confirmation problem changes character; it does not disappear. §25’s first limit is that the operator cannot know a command was received, and that is still true under the inversion: a receive-only platform cannot tell anyone what reached it. **What the inversion supplies is not confirmation of the platform’s present state but a bound on it.** The operator knows that the platform cannot remain enabled for longer than one heartbeat period after the last authorisation it actually received, whenever that was. That converts an uncertainty about safety into an uncertainty about availability—the commander does not know whether the emplacement is still live, and does know that if it is not being reached it is not still armed. The exchange is worth making and it should be described as an exchange.

It is cheaper. The heartbeat’s windows are at times the platform already knows, so the receiver need not sample continuously for an unpredictable command. *[Derived.]*

Heartbeat period	Message	Processing gain	Receive chain	Worst-case safing
1 h	200 s	60 dB	40.0 mWh/day	1.1 h
1 h	40 s	53 dB	8.0 mWh/day	1.0 h
6 h	200 s	60 dB	6.7 mWh/day	6.1 h
12 h	200 s	60 dB	3.3 mWh/day	12.1 h
24 h	200 s	60 dB	1.7 mWh/day	24.1 h

At a six-hourly heartbeat the receive chain costs **6.7 mWh a day**, which is less than half the corrected cost of recall-on-command, at full processing gain, with a bounded worst case that is stated rather than hoped for.

It supplies, in function though not in the instrument’s definition, the safeguard this architecture otherwise works against. §21.2 records an awkward result: Amended Protocol II’s *self-deactivation*—rendering a munition inoperable “by means of the irreversible exhaustion of a component, for example, a battery, that is essential to the operation of the munition”—is the one mechanism the paper’s whole energy argument is directed against, because a system engineered so that its battery never runs down cannot self-deactivate in that sense. **A heartbeat does not meet the instrument’s definition and this paper does not claim that it does: nothing is irreversibly exhausted, and the timer is renewed rather than spent.** What it is, and what it should be called, is an **automatically functioning loss-of-authorisation safing mechanism**: a countdown that expires unless a positive authenticated act from someone who is still present resets it. It occupies the functional position self-deactivation occupies—a safeguard that operates when nothing reaches the system—without the energy limit that the instrument’s mechanism depends on and that this architecture removes. Whether an instrument should recognise such a mechanism is a drafting question and not an engineering one; that a cheap mechanism exists in that position is the engineering finding.

And it gives paragraph 35(c) something to be tested against. The adopted text lists, among the measures of control, “measures to intervene in the operation of LAWS, to the extent

possible or feasible, including to deactivate or neutralize LAWS in a timely manner”. Under recall-on-command that is a promise about a message that may not arrive, whose worst case is four minutes when nothing is jamming and unbounded when something is. Under the inversion it is a bounded property of the design: **the system deactivates within one heartbeat period of the operator ceasing to authorise it, whether or not anything reaches it, and whether or not the operator is still alive to send it.** Whether any particular period is “timely” is a contextual judgement the text does not make and this paper does not make for it—six hours may be timely for a static obstacle in a rear area and plainly untimely where a friendly advance is expected within the hour. **What the inversion supplies is that there is an interval at all, that it is stated, and that it holds under denial.** That is the property the 18 December 2025 text’s dropped route had and the adopted text’s surviving route does not.

Three costs, stated because the inversion is not free.

Availability becomes the operator’s problem. A benign link outage safes a system the commander wanted armed. That is a real operational cost and it is the reason a designer would resist the inversion. It is also the correct direction for the cost to fall: the failure is an emplacement that stops working, borne by the force that emplaced it, rather than an emplacement that keeps working, borne by whoever walks into it.

The station must transmit on a schedule, which is an emission it would not otherwise make. §25.4 is the answer: the emission is at the station, to the rear, and can be placed on a remote or decoy antenna. But the schedule is now a pattern, and a pattern is itself a signature, so the transmission times should be drawn from the same cryptographic material as the listening schedule rather than being periodic in the clock.

Replay is now the attack that matters. An adversary who records one *continue* message and repeats it keeps the platform armed indefinitely. This is precisely what the monotonic counter and the authentication tag in §25.2 exist to prevent, and it is why the message cannot be shortened to the sixty-four bits an earlier version of this paper assumed. The inversion does not create the requirement for authenticated freshness; it makes it load-bearing.

The recommendation. Carry both. The heartbeat sets the outer bound on how long the platform may remain armed without positive authorisation and is what a timeliness obligation should be read against; the command channel remains for the faster and more specific actions—narrow the sector, lower the readiness level, safe it now—that a commander will want when the link is working. The heartbeat is what the system does when nobody is there. The command channel is what it does when somebody is.

25.4 Where the emission sits, and one fielded comparator

One fielded comparator is now available for the link this section leaves open, and it should be recorded for what it is worth and no more. The manufacturer of a seismic unattended ground sensor family states a core transmit power of **15 mW**, gateway-to-sensor standoff distances of **one to five kilometres with options to ten**, and describes the system as incorporating “low-probability of intercept and low-probability of detection” features.¹¹⁵ That is a buried, battery-powered node closing a link of the order this paper’s §15 design case requires, at a transmit power three orders of magnitude below a tactical radio. It does not close the link budget the Appendix calls for, because it states neither a jam margin nor a required sensitivity and because the geometry is a

¹¹⁵Agassounon and Steadman (n 84).

node transmitting rather than a node listening. What it does establish is that the range is not the implausible part.

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 **17.0 mWh per day on the command architecture and 6.7 on the heartbeat**—between one part in seventy and one part in a hundred and eighty 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.

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 anti-personnel 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.¹¹⁶ 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.

The instrument’s own vocabulary, set out at §21.2, should be used for both layers. The primary expiry this paper specifies is self-neutralization in the sense of article 2(11)—an incorporated mechanism that renders the platform inoperable—and self-destruction under article 2(10) is available on the same terms. The backstop the instrument pairs with them is self-deactivation under article 2(12), the irreversible exhaustion of a component essential to operation, and for this architecture that has to be a dedicated component whose exhaustion is irreversible and which sits on none of the harvested, exchanged or scavenged energy paths of Part VI, because Part VI is engineered to prevent exactly that exhaustion of the main reserve. The energy architecture creates that requirement; it does not remove it.

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.

¹¹⁶Amended Protocol II (n 107), Technical Annex.

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, and it is available to any system on this architecture. **What it costs is only partly known.** The worked receive costs are 17.0 mWh per day for the stated command-to-stop traffic assumption and 6.7 mWh per day for the six-hour permission-to-continue variant. **The energy and reliability of the independent physical expiry or irreversibly exhausting backstop have not been measured and must not be priced as zero:** a mechanical timer, a pyrotechnic link or a sacrificial cell has a mass, a shelf life, a temperature dependence and a failure rate, and this paper has measured none of them.

The pairing maps onto the text now before the review conference, and the mapping is worth stating. The report the Group of Governmental Experts adopted on 4 September 2026 lists at paragraph 35(c), among the measures by which control and human judgement are exercised, “measures to intervene in the operation of LAWS, to the extent possible or feasible, including to deactivate or neutralize LAWS in a timely manner”, which is a requirement about what a person can do to the system and is what the receive channel at §25 supplies. The text of 18 December 2025 had also offered, as an alternative route, that the system “incorporate self-destruct, self-deactivation or self-neutralization mechanisms”, which is what this section supplies and which the adopted text, like every 2026 draft before it, no longer names.¹¹⁷

27. 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.

¹¹⁷Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, Report of the 2024-2025-2026 sessions, UN Doc CCW/GGE.1/2026/3 (advance copy, 14 September 2026) paragraph 35(c); and rolling text, status 18 December 2025, section III paragraph 7(D). The Group adopted the report on 4 September 2026, as contained in CCW/GGE.1/2026/CRP.1/Rev.2 as orally amended; the advance copy is issued without formal editing, and the paragraph numbering should be checked against the edited document when it appears. In the adopted text the set of elements is paragraphs 25 to 44 of the report’s conclusions, numbered continuously; 35(c) is identified as the successor of III.7(D) by its wording and its place in the sequence of measures, and it differs from the draft of 3 September 2026 (CCW/GGE.1/2026/CRP.1/Rev.1, paragraph 37(e), “That LAWS can be deactivated or neutralized in a timely manner”) in framing the measure as one of intervention and qualifying it “to the extent possible or feasible”. The words “self-destruct”, “self-deactivation” and “self-neutralization” appear nowhere in the adopted text, and neither does “defined perimeter”; “incorporate” survives only in an unrelated sense (“those incorporating emerging technologies”, paragraph 27). The same words were absent from the drafts of 7 July and 3 September 2026 and from the rolling text of 5 June 2026. The companion papers cite the December text.

27.1 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, and how it fails depends entirely on which way the recall channel points.

There is no control link to sever. At posture 0 there is nothing. At posture 1 the outcome divides:

- **On a command-to-stop channel**, jamming denies recall, which harms the operator's ability to *stop* the system and does nothing to prevent it functioning. Stated plainly because it is counter-intuitive: **jamming such a system makes it more dangerous, not less**, and an adversary who understands the architecture has an active reason not to jam it.
- **On the permission-to-continue channel of §25.3**, the same jamming withholds the authorisation the platform needs to stay armed, and the platform safes itself within one heartbeat period. The adversary's cheapest attack becomes the safety mechanism.

That is not a difference of degree; it is a reversal of sign, and it is produced by the sense of one message. Everything else in this section applies to both, and the perverse-incentive result belongs only to the first.

The jam budget adds the constraint that makes this operational rather than rhetorical. Because the recall message is a few bytes that may be repeated for hours, its information rate—and therefore its processing gain—is a design variable, and at one bit per second against a megahertz of spread bandwidth the gain is 60 dB. On that figure a kilowatt jammer denies a five-kilometre recall link only within about **one kilometre of the platform**, and a ten-watt jammer within about three hundred metres. **To deny recall an adversary must therefore place an emitter within roughly a kilometre of an emplacement whose position they do not know, and leave it there.** That is a locatable, persistent commitment on the very ground the system exists to deny, and an adversary close enough to make it is close enough to be doing what §27.2 describes as the counter-measure that actually works.

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.

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.

27.2 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.**

27.3 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 §18.2 is not that it saves energy but that it removes a scheduled visit.

This also sets the correct target for the exchange model at §19. 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.

28. Unmitigated risks

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 seven 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 IX 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.

29. The denial test, and the safe-by-default emplacement

Part II supplies a test for closure: *does the mission fail if a party other than the operator withdraws a service?* That test asks what an adversary or a vendor can take away. It says nothing about what happens next, and this Part has repeatedly found that the second question is the one with the uncomfortable answers.

The test. For every safety-critical authorisation, constraint or assurance path, ask: *what state does the system enter when that path is unavailable, corrupted, stale or denied?* The safe response may be full safing, contraction of the usable envelope, or refusal of the affected function; it need not be identical for every service, and it is not required of paths that carry no safety-critical content.

The scope matters and an earlier version of this test overreached. A rule that every denied channel must stop the whole system is wrong: a reporting channel that fails deprives a commander of information without making a locally authorised system less safe, and requiring it to safe the platform would convert a nuisance into a mission failure for no safety gain. **The test bites on the paths that confer authority, enforce a constraint, or supply the evidence that a constraint still holds**—and on those it bites hard, because the natural way to build each of them fails it.

It is a one-line test and it is not a truism, because the natural way to build every channel in this specification fails it. A recall channel that carries a command to stop fails it: denying the channel keeps the system armed (§27.1). A geometric authority constraint fixed by survey fails it: displacing the platform silently relocates the constraint rather than breaking it (§15.4). A decision path with no record fails it: destroying or encrypting the compute module makes an autonomous engagement unprovable rather than impossible (§50). In each case the design behaves *better* when it is working and *worse* when it is attacked, which is the signature of a channel whose sense is inverted.

Applying the test to this specification yields three changes, all of which this paper has now made, and all of which are cheap. That they are cheap is the finding; that they are necessary is the argument.

Channel	Denied by	Naive design gives the adversary	Denial-safe design	Cost
Recall	jamming	an armed system that cannot be stopped	authenticated heartbeat: arm only while permission arrives (§25.3)	6.7 mWh/day
Emplacement geometry	loss of the survey-integrity premise through displacement	a system enforcing its sector over unsurveyed ground	inertial trigger plus visual landmark fit, ambiguity safes (§15.4)	microwatts, plus imagery already captured
Accountability	encrypting or destroying the compute module	an engagement that cannot be shown to have been autonomous	tamper-evident decision log, written to survive the platform (§50)	a few writes a day

One of the three is priced, one is nearly free, and one is not costed at all, and saying which is which matters more than an aggregate. The heartbeat is priced: 6.7 mWh per day at a six-hour period, *replacing* rather than adding to the 17.0 mWh recall line, so the denial-safe recall channel is cheaper than the one it replaces. The survey-integrity interlock adds an inertial channel in the microwatt class and runs its visual check on imagery tier three has already captured, so its marginal energy is small—though its *marginal false-safing rate* is not costed and is what A.5 measures. **The decision log is not costed and must not be reported as though it were.** A few non-volatile writes a day are cheap; surviving destruction of the platform, remaining readable

afterwards, protecting the record’s integrity against the party that holds it, and binding it to the configuration in force are four separate engineering problems, none of which this paper has specified and none of which it has priced. An earlier version of this section gave an aggregate figure for all three. That figure is withdrawn: **two of the three terms are quantified and the third is not.**

The point the aggregate was reaching for survives without it. Two of the three denial-safe properties cost less than the arrangements they replace, and none of the three is anywhere near the identification-and-release line that dominates the budget. That is the same result as §21.2’s, arrived at from the direction of failure rather than of capability: the properties that make an emplaced armed system tolerable are not the expensive parts of it. The expensive part is looking before acting, and everything that governs what happens when looking, listening or surveying is denied costs a rounding error on top.

Three observations follow, and the third is the one a rule-maker should take.

The test explains why the naive design is the one that gets built. Each inverted channel is the arrangement that behaves best in testing. A stop-command channel demonstrates beautifully on a range, where nothing is jamming. A surveyed sector demonstrates beautifully on ground nothing has displaced. An unlogged decision path demonstrates beautifully because nobody asks. The denial-safe alternative in each row is *worse* in a demonstration—it safes when the link drops, safes when the scene is obscured, and produces evidence that can be used against its operator. **The failure mode is systematically invisible at the moment the design is being judged,** which is a reason to require the property rather than to hope for it.

And it explains the shape of the treaty text. The 18 December 2025 draft offered two routes—deactivation on command, or an incorporated self-destruct, self-deactivation or self-neutralization mechanism. The first fails the denial test and the second passes it. The 2026 revisions, from the rolling text of 5 June to the report adopted on 4 September, kept the first and dropped the second (§51). On this analysis the drafting did not choose between two comparable safeguards. It kept the one an adversary can turn off and dropped the one an adversary cannot reach, and the arithmetic in this Part says the dropped one is also the cheaper.

The general form is worth stating because it outlives this specification. A rule written against a *capability*—what the system can do—is satisfied by a demonstration. A rule written against a *denial property*—what the system does when a named channel is taken away—is satisfied only by a design, because there is no way to demonstrate it without building it in. **Denial properties are therefore the class of requirement that a naming-based rule cannot express and a function-based rule can,** which is the argument of §40 arriving from the engineering side rather than the evidentiary one.

30. Doctrinal classification of the system

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**.¹¹⁸ 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.¹¹⁹ 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.

That property has already been broken in part, and by a programme of record rather than by a proposal. The M131 modular pack mine system permits one person to command a group of emplaced modules from several hundred metres, to destroy them and to recycle their self-destruct interval; the M7 Spider, fielded from around 2009, is a hand-emplaced networked munition under remote command whose components including undetonated munitions are recoverable; and the United States Army’s terrain-shaping obstacle effort is described in the Army’s own publications as carrying “on-off-on, remote destruct, self-locate and self-reporting capabilities” with networked two-way command and control, for the express purpose of recovering the effects of landmines while keeping them under friendly control after they are laid.¹²⁰ **An emplaced obstacle whose arming state can be changed from the rear without anyone going forward is therefore not new, and this paper does not claim it.** Part IX §47 sets the record out.

What is new is the axis on which the mutability runs, and stating it precisely is what makes the claim survive. Those systems toggle a munition of fixed lethal effect between armed and safe: the mutable quantity is *arming*, it is binary, and the effect the obstacle produces when armed was determined when it was built. A closed holding system at posture 1 is the first obstacle for which **the two components of obstacle intent that those systems leave fixed are themselves mutable after emplacement.** *Relative location* is fixed here as it is everywhere: 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, because the object carries a discriminating sensor whose behaviour is a parameter rather than a detonation state. The sector can be narrowed or widened within the arcs,

¹¹⁸US 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.

¹¹⁹On 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.

¹²⁰On the M131 modular pack mine system see Federation of American Scientists, ‘M131 Modular Pack Mine System (MOPMS)’ <https://man.fas.org/dod-101/sys/land/m131.htm> accessed 26 August 2026. On the M7 Spider see US Army, ‘Along came a spider’ (*army.mil*, 9 October 2012) https://www.army.mil/article/90105/along_came_a_spider/ accessed 26 August 2026. On the terrain-shaping obstacle effort see US Army, ‘Critical events in developing the concept of operations for terrain-shaping obstacles’ (*army.mil*, 2019) https://www.army.mil/article/224834/critical_events_in_developing_the_concept_of_operations_for_terrain_shaping_obstacles accessed 26 August 2026. Appendix A.7 records the withdrawal of the priority claim this passage previously made.

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 seventy 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, 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 Public Record

31. Statement of method

The sources for this compilation are published government budget documents, federal contract award records, granted patents and trade press. Published Comptroller exhibits print named programme elements and procurement lines for this capability class, each with its own identifier, title, budget activity and dollar figures for stated fiscal years. The federal award record lists named contracts, with contract numbers, performers, obligated amounts and periods of performance. Both statements are statements about published tables, and nothing more.

Classification is treated at the level at which the record speaks. In the FY2027 R-1 exhibit, every Army line item—1,337 of 1,337—carries the security-classification marking “U”. This is a statement about the classification column of that exhibit and about nothing else; it is not a statement about the R-2 narrative justification pages, about which this paper says nothing. The record also marks

its own boundary. The same exhibit carries rows whose Item column reads “Classified Programs”, with the placeholder programme element number 999999999 and the placeholder line number 999: dollar totals printed, no title, no narrative. And the justification volume’s own introduction states: “This document contains no classified data. Appropriately cleared individuals can obtain further information on Classified/Special Access Programs by contacting the Department of the Army.” Both passages are facts about the published document’s description of its own scope; neither is a statement about content, and the paper draws none.

Where the dated record stops, the stopping point is recorded as a date. For certain named efforts the last dated public entry can be identified, and the paper records that date as a fact about the record and nothing more: one programme’s last dated public entry is 26 September 2016; another effort’s public page states “This program is now complete”, with a timeline entry referring to 2015; for a third, the published address no longer resolves, checked 22 August 2026. No claim is made about what, if anything, followed. The absence of a public transition statement is not evidence of a non-public one. Where the secondary literature asserts that a status “is classified” without giving a source, the paper records that as a finding about the secondary literature and not about the status.

Where a document could not be reached by one route it was sought by another before being recorded as unavailable, and two conference papers first recorded as inaccessible were subsequently obtained and read in full. What remains unread is recorded as such at the point of citation. The consequence is a limit rather than a finding: the public record as compiled here does not permit any statement about what any unreached page contains.

Figures drawn from budget exhibits and award records are transcribed from the tables of the volumes cited, and each is given with the exhibit, date and line that carries it, so that any figure can be checked against its source.

This standard is stricter than the one the surrounding literature applies, and the paper adopts it because the whole value of a compilation from open sources lies in what it refuses to claim.

32. The funded public record

This section sets out what the published government documents and the federal contract award record contain for this capability class. It is drawn entirely from those sources. Each figure appears as printed in its source, with its own sentence; no figure is rounded, combined or compared, and no inference is drawn from the size or growth of any line about what any programme contains.

32.1 Budget exhibits

Exhibit R-1, Department of the Army FY2027 President’s Budget, Office of the Under Secretary of War (Comptroller), April 2026, prints programme element PE 0605053A, “Ground Robotics”, under Budget Activity 05 at line 143, with an FY2027 request of \$192.185m following an FY2026 enacted total of \$234.818m. The same element appears in the two preceding submissions with an FY2025 request of \$28.378m and an FY2026 request of \$227.038m. The same exhibit prints PE 0604017A, “Robotics Development”, under Budget Activity 04 at line 74, with an FY2027 request of \$20.290m. It prints PE 0604641A, “Tactical Unmanned Ground Vehicle (TUGV)”, under Budget Activity 05 at line 107, with an FY2025 actual of \$89.162m and no funding entry in either the FY2026 or the FY2027 column. It prints PE 0604531A, “Counter - Small Unmanned Aircraft Systems Advanced Development”, under Budget Activity 04 at line 96, carrying projects CQ5 and

CQ6, with figures of \$77.794m for FY2025, \$196.450m for FY2026 and \$43.300m for FY2027; the security-classification column for this line reads “u”.

Exhibit R-2, RDT&E Budget Item Justification, PB 2027 Army, April 2026, prints PE 0602135A, “Counter Small Unmanned Aerial Systems Applied Research”, project A31, with figures of \$0.000m for FY2025, \$12.618m for FY2026 and \$26.523m for FY2027.

Exhibit R-2, RDT&E Budget Item Justification, PB 2025 Army, March 2024, prints PE 0604601A, “Infantry Support Weapons”, project FF2, “Small Arms Fire Control”, with figures of \$7.880m for FY2023, \$10.050m for FY2024 and \$3.350m for FY2025. The same element and exhibit print project FM4, “Next Generation Squad Weapons”, with figures of \$17.156m for FY2023, \$16.141m for FY2024 and \$10.805m for FY2025. The same element and exhibit print project EW4, “Crew Served Weapons Engineering Development”, with figures of \$7.277m for FY2023, \$4.300m for FY2024 and \$3.685m for FY2025.

The R-1 gives no narrative description; it is a funding table only, and no statement about capability content is drawn from the size or growth of any line above. A blank cell in these exhibits is not, on its face, evidence of termination, transfer or classification, and no such claim is made about the FY2026 or FY2027 columns of PE 0604641A or of any other line.

32.2 Procurement lines

Exhibit P-40, PB 2027 Army, April 2026, prints P-1 line item 8310G04700, “Common Remotely Operated Weapons Station”, under Budget Activity 02 (Weapons) and Budget Sub-Activity 10 (Weapons). The Other Procurement, Army account carries line 92, “Robotics and Applique Systems”. The same account carries lines 57, 88 and 53 for ground sensing. The P-40 dollar and quantity table for the weapons-station line was not obtained; the line is identified here, and its figures are not given.

32.3 The federal award record

The obligated amounts and periods of performance below were read directly from the federal award data.

For remote weapon stations, contract W15QKN12C0103 records Kongsberg Defence & Aerospace AS, \$644,354,220.66, with a period of performance from 2012-08-17 to 2020-10-21. Contract W15QKN23F0074 records Kongsberg, \$51,382,044.00, from 2022-12-16 to 2026-02-28. Contract W15QKN23F0371 records Kongsberg, \$13,184,449.51, from 2023-06-29 to 2025-07-31.

For fire control as a separately contracted function, contract W15QKN23F0591 records Decilog Inc., \$1,230,109.63, from 2023-09-06 to 2026-09-05, with the description printed as “TASK ORDER AWARD FOR SMALL ARMS FIRE CONTROL AND OPTICS TECHNOLOGY REMOTE WEAPON STATION FIRE CONTROL ENGINEERING SUPPORT (SAFCOT RWSFC).”

For a weapon station combining automation and counter-UAS, contract W911NF25CA031 records Allen Control Systems Inc., \$1,551,553.53, from 2025-10-01 to 2026-11-30, with the description printed as “BULLFROG AI POWERED CUAS REMOTE WEAPON STATION FEASIBILITY STUDY”.

For aided target recognition, contract W5170124C0254 records Physical Sciences Inc., \$1,999,722.90, from 2024-09-12 to 2026-09-11, with the description printed as “SBIR PHASE II SEQUENTIAL TOPIC A20-043, INNOVATIVE APPROACHES FOR AIDED TARGET RECOGNITION (AITR)

OF ARMY TARGETS.” Contract W912CG26FA002 records the University of Southern California, \$211,322.00, from 2026-02-20 to 2026-12-17, “AIDED TARGET RECOGNITION (AITR) SCALE”. Contract W909MY21C0019 records Lynntech Inc., \$1,155,985.25, from 2021-04-28 to 2023-08-26, “DEEP GENERATIVE MODELING OF INFRARED DATASETS FOR AIDED TARGET RECOGNITION”.

For uncrewed ground systems in serial production, contract W56HZV23F0007 records General Dynamics Land Systems, \$43,578,127.03, from 2022-12-21 to 2026-09-14, for 135 full-rate production systems and ancillary kits. Contract W56HZV24F0036 records GDLS, \$31,586,922.76, from 2024-02-02 to 2026-09-30, for 145 systems. Contract W56HZV22F0004 records GDLS, \$21,341,026.31, from 2021-10-15 to 2026-08-14, for 144 low-rate initial production systems. These records describe an equipment-transport system; they do not on their face state any armament fit, and none is asserted here.

For an armed uncrewed ground system, contract M6785422F1003 records Oshkosh Defense LLC, \$22,210,813.37 obligated with outlays of \$14,846,913.88, from 2022-09-29 to 2025-03-31, “PRODUCTION OF THE REMOTELY OPERATED GROUND UNIT EXPEDITIONARY - FIRES (ROGUE-FIRES) CARRIER”. Contract M6785424F1007 records Hodges Transportation Inc., \$1,671,478.62, from 2024-04-08 to 2025-08-22, for test and evaluation including the “SEMI-AUTONOMOUS RESPONSE” of the ROGUE-Fires unit.

For crewed and uncrewed formations, contract W5170126CA118 records Primordial Labs Inc., \$1,899,925.03, from 2026-04-15 to 2027-04-15, “A244-P001 - VOICE COMMAND AUTONOMY FOR ARMY HUMAN-MACHINE INTEGRATED FORMATIONS”. Contract W912CH25FA121 records ASRC Federal Advanced Research LLC, \$400,049.97, from 2025-09-02 to 2026-12-31, for field service technician support for Human-Machine Integrated Formations.

For counter-UAS gun systems and ammunition, contract W519TC24F0158 records Northrop Grumman Systems Corporation, \$273,530,052.49, from 2024-05-14 to 2030-03-31, with the description printed as “LIGHTWEIGHT 30 XM1211 HIGH EXPLOSIVE PROXIMITY (HEP) CARTRIDGES. THE LIGHTWEIGHT (LW) 30MM XM1211 HIGH EXPLOSIVE PROXIMITY (HEP) URGENT MATERIEL RELEASE CARTRIDGE SUPPORTS GROUND-BASED PLATFORMS UTILIZING THE XM914 CHAIN GUN.” Contract W31P4Q24F0200 records SRC Inc., \$23,583,305.30, from 2024-09-03 to 2025-09-06, for six Increment 2.1 systems for one integrated-defeat system and ten for another. Contract W31P4Q23F0133 records SRC Inc., \$48,254,003.69, from 2023-07-19 to 2024-07-27, for fleet-wide sustainment and logistics support.

For photon-counting lidar, contract FA237724CB052 records Attollo Engineering LLC, \$749,879.90, from 2024-06-21 to 2026-09-21, “NOVEL, LOW C-SWAP, 2-MICRON GEIGER MODE APD ARRAYS FOR ATTRITABLE PLATFORMS”. Contract FA237724PB017 records Fibertek Inc., \$179,845.85, from 2024-06-13 to 2024-12-13, “COMPACT GEIGER MODE 3D LIDAR AND RANGE FINDER.” Contract FA865022C1023 records 3D-Sensir Inc., \$566,693.00, from 2022-06-01 to 2024-03-31, “ASYNCHRONOUS GEIGER MODE LADAR RECEIVER DEVELOPMENT”.

32.4 What the section establishes, and what it does not

Across the budget exhibits, the procurement lines and the award record set out above, these functions—uncrewed ground platforms, remote weapon stations, the fire control that aims them, target recognition, and photon-counting optical sensing—are separately named, separately funded and separately contracted, with entries current to 2026. They are not a projection. The section

does not establish what any of these systems can do, how well any of them works, or whether any of them has been fielded in the configuration this paper specifies. The award record states performer, amount and period of performance; it does not state what was delivered.

33. Ballistic computation as a catalogue item and a funded line

Every specification figure in the paragraphs that follow is a manufacturer's sales figure, taken from published catalogue or marketing material, not an independent measurement; the manufacturers' own wording is reproduced because it is the evidence, and the paper records what is printed without drawing any engineering inference beyond it.

One manufacturer advertises "AUTOMATIC TARGET TRACKING" and "AUTOMATIC BALLISTIC CALCULATION" as headline features of a remote weapon station family; the spelling of "BALLISTIC" is the manufacturer's and is reproduced here as printed. A second manufacturer publicly specifies a 47 kg remotely operated weapon module which, on a single "fire" command, autonomously calculates all impact points, slews the weapon and fires, and the same page specifies "map-based targeting through the turret control software with an integrated ballistic calculator" accounting for "weather, grenade temperature, ammunition type, height differences". The "autonomously calculates all impact points... and fires" wording describes a multi-round automatic sequence initiated by one human command, and is recorded as printed with no inference beyond that. A third manufacturer advertises a station with a "sophisticated ballistic solution" for first-round hit probability. A fourth advertises a medium station with a "laser range finder for ballistic calculation". A US Army Center for Army Lessons Learned paper describes a fielded air-defence turret as performing "ballistic calculation. Adjust for speed (up to 800 kilometers per hour), wind, and predicting the precise point of interception"; that paper's own footnotes 1, 4 and 6 all cite a single trade article of 27 May 2025.¹²¹

The same function appears as a separately funded government project line. Programme element 0604601A, "Infantry Support Weapons", project FF2, "Small Arms Fire Control", appears in Exhibit R-2 of the Army PB 2025 justification, dated March 2024, with the figures set out at §32.1. Fire control is therefore a separately identified and separately funded project line and not only a catalogue feature; the R-2 gives figures and a project title, and the paper says nothing about what the work comprises. Two granted patents are recorded here for existence and title only: US 12,259,460 B2, "Camera systems and devices for ballistic parameter measurements", granted 25 March 2025; and US 10,803,758 B1, "Wind prediction systems and methods", Lockheed Martin Corporation, inventor Keith S. Barr, priority 21 December 2016, granted 13 October 2020. The paper has not read the claims of either patent and says nothing about what they cover.

One nil return should be recorded. No source in this compilation names earth-curvature, Coriolis or spin-drift correction as a module; the nearest is one programme's stated goal of measuring "all relevant physical phenomena that influence the ballistic trajectory", which names no phenomenon in particular. The proposition these paragraphs establish is only this: that a module which ingests environmental variables and computes a firing solution is a line item in a manufacturer's catalogue and a funded government project line, and is not a research aspiration. The paragraphs make no comparison among the products named, rank none of them, and say nothing about what any of them could be made to do.

¹²¹Center for Army Lessons Learned, US Army Combined Arms Center, No. 26-1140 (March 2026), whose footnotes 1, 4 and 6 cite a TechUkraine article of 27 May 2025 that itself credits the United24 Media article of 26 May 2025 from which the Sky Sentinel reporting at §52 derives.

34. What is on sale to the public

The preceding sections record what governments fund and what defence manufacturers advertise. This section records what is offered to the general public, because the three functions the architecture at Part V combines—a platform that carries, a sensor that perceives and an effector—are each listed at consumer prices, and the listings are the evidence. Every figure is as printed on the vendor’s own page on the date given, and the paper draws no inference about what any product could be made to do.

A general-purpose legged platform. A manufacturer’s product page lists a humanoid platform at “US \$13.5K” in its standard configuration, excluding taxes and shipping, with a research configuration priced on application.¹²² The same manufacturer publishes, at a public code repository, a software development kit for its platforms in C++ and Python, ROS 2 packages for two of its quadrupeds, and reinforcement-learning training environments for its legged platforms.¹²³ The paper records the listing and the repository; it has not evaluated either.

One further fact about the same manufacturer belongs here, and it is recorded because it cuts against this paper’s own use of the product rather than in spite of it. That manufacturer is among the signatories to an open letter of October 2022, published with several other makers of general-purpose robots, undertaking not to weaponise its platforms.¹²⁴ The signatories also state that because of the permissive licences their software carries, it is “difficult, if not impossible” for them to limit others from using it to build weaponised systems, and that they do not condone such efforts. A stated policy against weaponisation, resting on a conceded inability to prevent it, is the clearest illustration the record offers of the point §21.3 draws: the restraint is a policy layer sitting on a commodity, and the commodity is supplied on terms its maker cannot enforce.

An armed commercial product. A second manufacturer offers for sale, at USD 9,420, a flame-projecting device mounted on a commercial quadruped robot, described on its own page as “the world’s first flamethrower-wielding robot dog”, with lidar mapping for terrain awareness and obstacle detection, laser sighting, first-person video navigation from a handheld controller, a 30-foot flame range and a 350-foot line-of-sight operating range; the page states that the product is “legal in 48 U.S. states” and that its flamethrower component is federally unregulated and not considered a firearm.¹²⁵ The effect is incendiary rather than ballistic and the device is remotely operated. The paper records only that a weapon mounted on a commercial robot has been offered to the public at a four-figure price, within the law as the vendor states it.

¹²²Unitree Robotics, ‘Humanoid robot G1’, product page <https://www.unitree.com/g1> accessed 24 August 2026, listing the standard configuration at “US \$13.5K” excluding taxes and shipping and the EDU configuration as “Contact sales”. A manufacturer’s listing.

¹²³Unitree Robotics, public repositories at <https://github.com/unitreerobotics> accessed 24 August 2026: `unitree_sdk2` (C++), `unitree_sdk2_python`, `unitree_ros2`, and the `unitree_rl_gym`, `unitree_rl_lab` and `unitree_rl_mjlab` training environments. Licence terms were not examined.

¹²⁴‘General Purpose Robots Should Not Be Weaponized: An Open Letter to the Robotics Industry and our Communities’ (6 October 2022) <https://www.openrobotics.org/blog/2022/10/6/general-purpose-robots-should-not-be-weaponized-an-open-letter-to-the-robotics-industry-and-our-communities> accessed 26 August 2026. The signatories include the manufacturer whose product page is cited at (n 122). The paper records the undertaking and takes no view on it.

¹²⁵Throwflame, ‘Thermonator’, product page <https://throwflame.com/products/thermonator-robodog/> accessed 24 August 2026. The price, the feature list, the statement that the product is “legal in 48 U.S. states” with restrictions in California and Maryland, and the characterisation of the flamethrower component as federally unregulated and not a firearm are the vendor’s. Coverage at launch: Gizmodo, ‘Thermonator, the Flame-Throwing Robot Dog, Can Now Be Yours for \$9,420’ (2024).

An armed platform shown to the trade. At a trade exposition in October 2021 a quadruped platform was displayed fitted with a 6.5 mm rifle payload; a product card circulating at the show stated “precision fire out to 1,200 meters”, and the platform manufacturer subsequently stated that it had not produced the card and that the system “is entirely operator-controlled and does not have any autonomous capabilities”.¹²⁶ Its chief executive stated the same month that the payload “has no autonomy and no AI”, that “it’s triggered from a distance, and that has to be done by a human”, and that “there is always an operator in the loop”.¹²⁷ The paper takes those statements at face value. What they describe is a decision about the trigger. Everything upstream of it—mobility, stabilisation, the sight—is stated to work in the same way whichever way the decision is taken, and §15.7 returns to that point.

Open middleware. The Robot Operating System in its second generation is the open-source software substrate on which this class of platform is ordinarily developed, and the manufacturer above publishes packages for it;¹²⁸ lidar of the kind the two products above carry is a commodity sensor; and the operating principles of ranging and detection are the open literature set out at §35. None of that is a finding. It is the reason the listings above are the relevant fact.

What the section establishes is narrow. A platform, a perception stack and an effector are each on sale to the public, separately and lawfully as the vendors state it. The section does not establish that any purchaser has combined them, and it says nothing about what a combination could do. The consequence for this paper is the one §21.3 draws: the components that are commodities are the platform, the sensors and the effector, and what no listing offers is the judgement that the 91 per cent of the budget at §21.2 buys.

35. The sensing principles, from the open literature

This section explains how each class of sensing and computation the paper relies on works in principle. It is drawn entirely from open peer-reviewed literature and government technical reports, it reproduces no equation, and it applies no relation to any particular system. The purpose is to show that the operating principles are textbook and the limits calculable from published relations, so that a reader can reason about what is and is not physically possible without any non-public source. Where a source is marked below as not directly read, the paper cites it for the stated proposition rather than asserting the proposition independently.

35.1 Coherent detection, and why aperture does not solve everything

A heterodyne optical receiver is at once a receiver and an antenna, and its effective aperture and field-of-view solid angle cannot be chosen independently: their product is fixed by the wavelength, as established by A. E. Siegman, “The antenna properties of optical heterodyne receivers”, *Proceedings of the IEEE* 54(10), 1350–1356 (1966), with a twin paper at *Applied Optics* 5(10), 1588–1594 (1966). For a diffuse target, W. R. Leeb, P. J. Winzer and K. H. Kudiella, *Applied*

¹²⁶The War Zone, ‘Robot Dogs Now Have Assault Rifles Mounted On Their Backs’ (12 October 2021, updated 15 October 2021) <https://www.twz.com/42717/robot-dogs-can-now-have-6-5mm-assault-rifles-mounted-on-their-backs> accessed 24 August 2026, reporting the display at the Association of the United States Army convention of a Ghost Robotics quadruped carrying SWORD International’s Special Purpose Unmanned Rifle in 6.5 mm Creedmoor, the product card’s “precision fire out to 1,200 meters”, and the manufacturer’s subsequent statement that the card was not its own and that the system is entirely operator-controlled.

¹²⁷Evan Ackerman, ‘Q&A: Ghost Robotics CEO on Armed Robots for the U.S. Military’ (*IEEE Spectrum*, 18 October 2021) <https://spectrum.ieee.org/ghost-robotics-armed-military-robots> accessed 24 August 2026.

¹²⁸Open Robotics, ROS 2 documentation <https://docs.ros.org> accessed 24 August 2026; and the `unitree_ros2` package at (n 123).

Optics 37(15), 3143–3148 (1998), show that mixing efficiency falls as the receive aperture grows, so that signal-to-noise saturates rather than scaling with aperture area. R. G. Frehlich and M. J. Kavaya, *Applied Optics* 30(36), 5325–5352 (1991), cited here and not read, is cited for the proposition that turbulence degrades mixing efficiency further and that performance must be computed with turbulence included.

35.2 Frequency-modulated ranging

With a triangular chirp, range maps to the mean of the two beat frequencies and radial velocity to their separation, as set out by J. Riemensberger and others, *Nature* 581, 164–170 (2020) (arXiv:1912.11374). D. Uttam and B. Culshaw, *Journal of Lightwave Technology* 3(5), 971–977 (1985), cited and not read, is cited for the proposition that range resolution is set by the chirp excursion bandwidth and maximum range by source coherence length. P. A. Roos and others, *Optics Letters* 34(23), 3692–3694 (2009), cited and not read, is cited for the proposition that chirp non-linearity, rather than laser linewidth alone, is the dominant limit on range precision. Q. Na and others, *Optics and Lasers in Engineering* 164, art. 107523 (2023), cited and not read, is cited for the proposition that a fixed optical frequency offset on one arm removes the sign ambiguity in the Doppler term.

35.3 Speckle, and why it sets a floor

Fully developed speckle from a rough surface has negative-exponential intensity statistics—the standard deviation of the returned intensity equals its mean—and its contrast falls only as the inverse square root of the number of independent realisations averaged on an intensity basis, as established by J. W. Goodman, *Journal of the Optical Society of America* 66(11), 1145–1150 (1976), and *Speckle Phenomena in Optics*, 2nd ed. (SPIE Press, 2020). N. Bender and others, *Optica* 5(5), 595–600 (2018) (arXiv:1711.11128), show that those statistics are the default outcome of a fully developed circular-Gaussian field rather than a law of nature.

35.4 Photon-counting detection

An avalanche photodiode biased above breakdown detects single photons, and its performance is characterised by photon detection efficiency, dark count rate, afterpulsing and timing jitter; S. Cova and others, *Applied Optics* 35(12), 1956–1976 (1996), **now read in full**, is cited for that characterisation; earlier versions of this compilation recorded it as cited and not read, and the correction is recorded at Appendix A.7. J. Rapp and others, *IEEE Signal Processing Magazine* 37(4), 62–71 (2020), establish that single-photon data are photon counts rather than samples of a linear system with Gaussian noise, so that the estimation problem must be posed in Poisson statistics. D. Shin and others, *IEEE Transactions on Computational Imaging* 1(2), 112–125 (2015) (arXiv:1406.1761), show that depth and reflectivity can be estimated from on the order of one detected photon per pixel when the estimator uses the exact counting statistics. A. W. Ziarkash and others, *Scientific Reports* 8, art. 5076 (2018), report that standard afterpulsing models disagree with each other and with measurement.

35.5 Propagation through the atmosphere

Transmittance follows Beer–Lambert extinction, and the extinction coefficient can be related to measured visibility, as given by I. I. Kim, B. McArthur and E. Korevaar, *Proc. SPIE* 4214, 26–37 (2001).

The negative result that source carries on wavelength, and the turbulence framework from which scintillation, beam wander and coherence length follow, are set out at §36, which states the conditions of use.

35.6 What is physically achievable, independent of hardware

The variance of any unbiased estimator of delay, Doppler or angle is bounded below by the inverse Fisher information; the canonical statements are H. L. Van Trees, *Detection, Estimation, and Modulation Theory*, Parts I and III (Wiley, 1968 and 1971), and S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I* (Prentice Hall, 1993), both cited and not read. Applied to accumulated heterodyne spectra, B. J. Rye and R. M. Hardesty, *IEEE Transactions on Geoscience and Remote Sensing* 31(1), 16–27 (1993), show that this bound sets the floor on velocity precision. The corresponding bound for direction-of-arrival estimation, and the conditions under which subspace methods attain it, are given by P. Stoica and A. Nehorai, *IEEE Trans. ASSP* 37(5), 720–741 (1989), cited and not read.

This heading matters most to the paper’s argument, and it is worth saying why. A bound of this kind is a statement about what no hardware can beat. It is available to a rule-maker without any knowledge of any particular system, and it is the reason a rule can be written against function rather than against a device.

35.7 Wind, and the computation of a trajectory

The crosswind component transverse to an optical path can be recovered from the drift of the scintillation pattern, as shown by T. Wang, G. R. Ochs and R. S. Lawrence, *Applied Optics* 20(23), 4073–4081 (1981). O. Porat and J. Shapira, *Applied Optics* 49(28), 5236–5244 (2010), cited and not read, is cited for the proposition that crosswind can also be estimated passively from turbulence-induced image-motion statistics recorded by an ordinary video camera. Z. Liu and others, *Remote Sensing* 11(21), art. 2522 (2019), describe pulsed Doppler wind lidar measuring radial velocity from the Doppler shift of aerosol backscatter, with range resolution set by pulse width and range gating. Trajectory computation itself is open and standardised as a modified point-mass model, which augments point-mass translation with an estimated yaw of repose and an explicit axial spin equation; the sources are R. F. Lieske and M. L. Reiter, Ballistic Research Laboratories Report 1314, US Army Materiel Command, March 1966 (DTIC AD0485869), L. Baranowski, P. Majewski and J. Szymonik, *Bulletin of the Polish Academy of Sciences: Technical Sciences* 68(5), 1167–1175 (2020), and NATO STANAG 4355, “The Modified Point Mass and Five Degrees of Freedom Trajectory Models”, cited and not read. R. L. McCoy, Technical Report ARBRL-TR-02293, US Army Ballistic Research Laboratory, February 1981 (DTIC ADA098110), establishes that zero-yaw drag can be estimated from projectile geometry alone.

Every relation named above is published, most of it decades old, and none of it is in dispute. The physics that determines what an optical sensing and fire-control architecture can do is open literature. What is not public is which of these arrangements a given party has built and how well it works—and a rule drafted against that is a rule drafted against the one thing nobody outside can check. A rule drafted against function does not have that problem, because function is calculable from the relations above by anyone who can read them.

36. Atmospheric propagation: conditions and limits

The established framework for optical attenuation in the atmosphere is set out in a NASA Jet Propulsion Laboratory monograph chapter (Biswas and Piazzolla, *Deep Space Communications*, chapter 3, 2008). It describes attenuation along a path by the Beer–Lambert relation, and decomposes the total attenuation coefficient into molecular absorption, molecular scattering, aerosol absorption and aerosol scattering; it gives the aerosol-scattering term as a function of visual range and of wavelength through a power law. These relationships are named here and not reproduced, because the purpose of the passage is to establish conditions of use rather than to carry the arithmetic; the source states them in full.

The single most important point in this passage is a negative one, and it corrects a claim that is widely repeated in the field. It is often said that operating at around 1.5 micrometres buys penetration through fog and poor weather. The source most often cited for the general propagation model is Kim, McArthur and Korevaar, *Proceedings of SPIE* volume 4214 (2001), which gives a visibility-based attenuation model in which the particle size-distribution exponent is an explicit function of visibility. That paper has been read in full for this compilation, and it is worth quoting because it corrects the misconception explicitly and supplies the mechanism. The wavelength dependence in the model is carried entirely by the size-distribution exponent, and the paper’s finding is that in fog that exponent goes to zero, which removes the wavelength term altogether: its conclusion is that “785 nm, 850 nm, and 1550 nm are all attenuated equally by fog”, and that on the corrected expression “there is no longer any difference in attenuation losses transmitting at 1550 nm or 785 nm in fog conditions” below 500 m visibility. The advantage is real in haze, where the paper states that at visibilities above 2 km the prediction of less attenuation at 1550 nm “is most likely true”, and absent in fog. The magnitude of what is at stake is also given: attenuation over the range of conditions the paper treats runs from about 2 dB per kilometre in exceptionally clear weather to about 310 dB per kilometre in very dense fog. On that source, then, the longer wavelength buys eye safety; it does not buy weather penetration. **The correction comes from the paper usually cited for the opposite conclusion**, and the reason it is so often cited for the opposite conclusion is that the earlier expression it corrects is the one still in circulation.

Measured particle size distributions and extinction coefficients for battlefield obscurants exist in the open peer-reviewed literature across 0.249 to 12 micrometres (Gao and others, *Remote Sensing* volume 15 number 20, article 5038, 2023). The authors of that paper are at the Beijing Institute of Technology, which is not a neutral civil source. No numeric extinction coefficient at 1.5 micrometres specifically could be extracted from it, and nothing here should be read as supplying one; the dynamic range quoted above is from the free-space-optics source and concerns fog and haze rather than battlefield obscurants, and the two should not be conflated. On optical turbulence, the canonical source is L. C. Andrews and R. L. Phillips, *Laser Beam Propagation through Random Media*, 2nd ed., SPIE Press Monograph PM152 (2005), cited here as the canonical source and not read, which characterises turbulence by the refractive-index structure parameter and the Kolmogorov spectrum and derives scintillation, beam wander and coherence length from them; the two-thirds law is given in A. N. Kolmogorov, *Proceedings of the Royal Society of London A* 434, 9–13 (1991, English reprint of the 1941 paper); the coherence length that caps achievable angular resolution is given in D. L. Fried, *JOSA* 56(10), 1372–1379 (1966), cited and not read; and beam wander specifically is treated in M. A. Cox and others, arXiv:1907.10519 (2019), which states the scaling for a collimated Gaussian beam in Kolmogorov turbulence; the expression is not reproduced here. A published height-stratified structure-parameter profile also exists (Spiers, NASA-CR-196510, 1994); the profile is cited here as a published figure, its table having been extracted rather than read from

the original, without asserting a value at any particular height and without quoting any specific number from it.

One consequence follows from the framework rather than from any measurement of a particular system. Because aerosol scattering behaves as the framework above describes, a beam path is visible off-axis to a sensitive short-wave infrared imager in fog, dust, rain or smoke. That is a limit on any architecture that relies on an optical emission remaining private, and it follows from the propagation relationships named above and not from any system test.

What all of this establishes is not that a given system will or will not work, but that the conditions under which an optical sensing architecture is usable are published, bounded and checkable—which is what a rule-maker needs and what the vendor literature does not supply.

37. Two programmes in precision small-arms fire, and where their public account ends

This section records two government research programmes in precision small-arms fire, sets out what their own published sources state, and marks the date on which the public account of each ends. Every figure below is reported as its source reports it.

Programme one—crosswind measurement and ballistic computation. A programme was initiated to develop a system that profiles and compensates for down-range crosswind in order to improve first-round hit probability, with an initial Phase 1 contract of approximately \$2m. A subsequent contract of \$9.7m over 18 months covered integration of crosswind measurement into a prototype spotter scope that automatically transmits crosswind information to modify the crosshairs and display bullet strike location; the source’s headline gives \$3.9m for that contract while its text gives \$9.7m, and these funding figures are self-reported by the contractor. A \$6m contract dated 1 August 2012 covered a successor described as “a compact observation, measurement, and ballistic calculation system mountable to combat rifles or spotting scopes”. The stated goals recorded for the successor included measuring “all relevant physical phenomena that influence the ballistic trajectory”, rapid calculation and display of an offset aim point and a confidence metric in the shooter’s riflescope, operation from -20 to 49 degrees Celsius, and crosswind and range measurement within 3 seconds. This specification language is single-sourced to trade press, is unverified against a government primary document, and the article carrying it names no contractor. One element of that specification is the legally interesting part: the specification called for the device to output a measure of its own uncertainty alongside its solution.

The public account of programme one stops short of its method. No accessible public source states the crosswind sensing method used by either effort; across every source located, the measurement subsystem is described only functionally. No accessible source identifies the detection physics, the wavelength, the aperture, the transmitted power or the architecture. No live programme page exists for the successor, and its published address no longer resolves, checked 22 August 2026. The paper does not know the method, and no accessible source states it; the paper neither names a method nor reasons toward one from any adjacent technology.

Programme two—guided small-calibre projectile. Government publications record that the programme sought “to revolutionize rifle accuracy and range by developing the first ever guided small-caliber bullet”, a manoeuvrable .50-calibre round with real-time guidance which could “change path during flight”. The agency announced in April 2015 that the programme had demonstrated a self-steering .50-calibre round with real-time optical guidance, that an experienced shooter

“repeatedly hit moving and evading targets”, and that a novice using the system for the first time “hit a moving target”. A contractor states that it “recently demonstrated the first-ever precision guided munition capability in live fire testing of guided .50 caliber bullets”, with a team of four named companies and laboratories. A separate national laboratory publishes its own research page describing a .50-calibre self-guided projectile for a smooth-bore barrel. These publications are a government’s and a contractor’s self-report of the programme; no quantified hit rates, ranges, target speeds, sample sizes or conditions are published, and the paper reports what the sources state and no more.

The public account of programme two goes quiet after 2015, and the agency states that the programme ended that year. No government source states any transition to a programme of record, any service adoption, or any successor. The absence of a public transition statement is not evidence of a non-public one.

A dated instance of the inference this paper refuses. A trade article of 25 July 2026 asserts that the “precise ongoing military operational status is classified” and that “any production deployment is rarely discussed in public”; the article cites no source for the classification claim and reports no information post-dating the 2015 record. This is recorded here as a finding about the secondary literature and not about the programme: a dated, current instance, on this paper’s own subject matter, of exactly the inference this paper declines to draw.

The point of closing on that article is not that the trade press is careless. It is that a documented silence gets filled, that the filling then circulates, and that an instrument drafted against the circulating account inherits the invention. That is the argument for drafting against function rather than against programmes, and it is the cleanest illustration of it the paper has.

38. What the published record says about its own scope

This section records a narrow point: that the published government documents carry, on their own face, markers of material held outside them. It argues nothing about what any such material contains, and every sentence below can be checked by opening the document named and looking at the line, note, column or field it refers to.

Exhibit R-1, Department of the Army FY2027 President’s Budget, Appropriation 2040A RDT&E Army, April 2026, carries rows whose Item column reads “Classified Programs”, with the placeholder programme element number 999999999 and the placeholder line number 999. Dollar totals are printed for these rows, and the volume carries no titles, projects or narrative justification for them. The Budget Activity 02 row reads, in thousands of dollars, 35,766 / 34,599 / 34,599 / 32,883 / 32,883. The Budget Activity 04 row reads, in thousands of dollars, 10,262 / 7,764 / 7,764 / 22,889 / 22,889. The columns, as printed, are Line, Program Element, HQ Item, Sec, Act, FY2025 Actuals, FY2026 Discretionary Enacted, FY2026 PL 119-21 Spend Plan, FY2026 Total, FY2027 Discretionary Request, FY2027 Mandatory Request, FY2027 Total. The “Sec” column, which is the security-classification column, reads “u” for both rows. These rows are an accounting aggregate and not a named programme, and the paper says nothing whatever about what the aggregate comprises.

The introduction to the Army FY2027 RDT&E justification volume carries a note, which appears verbatim in the FY2026 and FY2025 volumes as well: “This document contains no classified data. Appropriately cleared individuals can obtain further information on Classified/Special Access

Programs by contacting the Department of the Army.” This note is volume-level front matter and attaches to no programme element.

The *DARPA Guide to Broad Agency Announcements and Research Announcements* provides that “If the nature of the BAA effort is classified or is anticipated to involve access to or generation of classified information, a DD Form 254, ‘DoD Contract Security Classification Specification,’ will be required as an attachment to the BAA,” and, in describing the special-considerations section, refers to the “existence of any classified addendums”. This document is general procedural guidance and not a specific solicitation; it establishes that the category exists, and it does not establish that any particular solicitation in this paper’s subject areas carries such an addendum, of which the paper asserts none. The document’s own date could not be reconciled with its posting path, and no date is stated for it here.

The Department of the Army’s component instructions to the DoD 24.4 Small Business Innovation Research annual announcement, posted 4 April 2024, provide at one topic that “The technology within this topic is restricted under the International Traffic in Arms Regulation (ITAR), 22 CFR Parts 120-130, which controls the export and import of defense-related material and services, including export of sensitive technical data, or the Export Administration Regulation (EAR), 15 CFR Parts 730-774, which controls dual use items,” and that “Successful firms will be required to comply with Controlled Unclassified Information (CUI) DoDI 5200.48.” This provision is a restriction on the handling and export of technical data, imposed by regulation, and it is a different thing from classification; the section marks the distinction rather than blurring it.

Within the federal award record, the description fields of sibling awards under the same system in the same fiscal year vary in length. Two 2024 task orders carry the description “M-LIDS TO 14” in full: W31P4Q24F0204, \$29,765,392.00, period of performance 2024-07-05 to 2026-09-30, and W31P4Q24F0221, \$27,002,472.75, period of performance 2024-07-25 to 2025-09-30. A record under the same system reads in full: “THE PURPOSE OF THIS TASK ORDER W31P4Q24F0200 IS TO PROVIDE SIX (6) INCREMENT (INC.) 2.1 SYSTEMS FOR FS-LIDS & TEN (10) INC. 2.1 SYSTEMS FOR M-LIDS.” Award W31P4Q25F0206, \$325,291,342.98, carries the description “COYOTE INTERCEPTORS” and nothing further. Brevity in an award description has many ordinary causes—data-entry practice, reference to a parent contract, administrative convention—and this section records the variation as a property of the records and offers no explanation for it.

What all of this establishes is narrow: the published record marks its own boundaries, and a reader of it knows where the published account stops. That is a fact about documents. The consequence for this paper’s argument is the only reason the section exists. An instrument drafted against what the public record describes is drafted against a record whose own front matter states that it is not the whole of the account. That is an argument for drafting against function and against use, rather than against named programmes and named technologies, because the naming is where the record is uneven.

39. The boundary of the record moves, and which way a rule-maker should err

§38 recorded that the published documents mark their own edges. This section records something further and more consequential for an instrument being drafted now: **those edges move, they have moved recently, and they moved in the direction of less disclosure while this paper was being written.** It then states the only conclusion this paper thinks can honestly

be drawn from that, which is a conclusion about how to draft under uncertainty rather than a conclusion about what is behind the edge.

39.1 The rule this section applies to withheld material

The rule is stated first because it is what makes the rest usable, and because the opposite practice is common enough in this field that a reader is entitled to know which one is in operation.

Where material has been withheld, redacted or removed, this paper records **that** it was withheld, **when, by whom**, and **what the withholding authority itself stated**. It draws nothing from the content of what is withheld, attempts no reconstruction of it, and treats the extent of withholding as evidence about the record and not as evidence about any capability.

The prohibition in the second sentence is not fastidiousness and it is not modesty. It is the same rule §31 imposes on the rest of Part VIII, and §37 closes on a dated instance of what happens when it is abandoned: a trade article asserting that a programme’s status “is classified”, citing nothing, on this paper’s own subject matter. An argument that reasons from a redaction to a capability is that article’s argument. It is also, and this is the practical point, an argument that cannot be checked, and an unfalsifiable claim put before a treaty body is worth less than no claim at all because it discredits the checkable ones beside it.

There is a further reason, which §38 already gave in a smaller case. Withholding has many ordinary causes—sources and methods, contractor proprietary data, export-control regulation, administrative convention, the handling categories §38 distinguishes from classification. §38 declined to explain the brevity of an award description and this section declines, for the same reason, to explain a redaction. What can be said is what a document records about itself, and it turns out that this is a great deal.

39.2 What the record says about its own contraction

The aggregate is not a fixed quantity. The rows §38 sets out print, for the same accounting aggregate, a series across fiscal years, and the series is not flat: the Budget Activity 04 row reads 10,262, then 7,764, then 22,889, in thousands of dollars. That is a statement about a printed row and nothing else; no inference is drawn here about what the aggregate comprises, and none can be. What it establishes is only that the quantity held outside the described record varies from year to year, so that a reader who takes a single year’s exhibit as the shape of the record has taken a snapshot of something that moves.

And the described record itself has recently contracted, on this paper’s own subject matter. The Office of the Director, Operational Test and Evaluation is the United States authority whose reports assess whether weapon systems work; its annual reports have been public for decades and are the source of some of the most specific public statements that exist about what particular systems do. Three dated facts are recorded here and nothing is added to them.

In late July 2026 the Department of Defense removed the landing pages through which twenty-five years of those annual reports had been reached, and reporting on the removal states that the documents themselves were not deleted and that many remain reachable by direct address.¹²⁹ The

¹²⁹Stavroula Pabst, ‘Why would the Pentagon want to hide weapons testing results?’ (*Responsible Statecraft*, 20 August 2026) <https://responsiblestatecraft.org/dod-weapons-testing/> accessed 26 August 2026, reporting that twenty-five years of annual reports of the Director, Operational Test and Evaluation were taken offline in the pre-

Department's stated reason, as reported, was that adversaries could exploit artificial intelligence "and other tools to compile disparate, unclassified data points" from the reports "into highly detailed vulnerability profiles".¹³⁰ In August 2026 two members of Congress wrote to the Secretary of Defense citing an audit of 30 June 2026 by the Government Accountability Office, which found that a reorganisation had reduced the office's capacity to oversee the 173 weapon system programmes on its 2025 oversight list and increased "the risk of weapon systems being delivered to the warfighters with undocumented shortfalls"; the letter records a staffing reduction of approximately seventy-four per cent and a budgetary reduction of almost eighty per cent, raises "the recent removal of DOT&E's annual reports from its official website", and requests a response by 26 August 2026.¹³¹ An earlier letter, of February 2022, had recorded that twenty-two programmes were redacted from that year's report and one removed from it entirely, following a decision that only a controlled-unclassified version would be provided.¹³²

Two observations follow, and both are observations about documents.

The first is that this paper is a specimen of what the change affects. The most useful single sentence in Part IX—a government test authority recording that a fielded munition "no longer has the capability to engage targets autonomously", which is this paper's evidence that an autonomous engagement function is separable and therefore regulable—is from that class of document.¹³³ The paper holds it because it was retrieved and cited by direct address. A reader coming to the question fresh, and browsing rather than holding the address, would now have a materially harder time finding it. The paper says only that: the documents this paper relies on were reached, the route by which a reader would discover them was removed, and the difference between those two statements is exactly the difference between an archive and a discoverable archive.

The second is that the stated reason and this paper's argument describe the same phenomenon and part company on what to do about it. The rationale reported for the removal is that unclassified data points can be aggregated into something more sensitive than any of them. Part VIII is an exercise in aggregating unclassified data points, and §41 states plainly

ceding month, and that the removal was of the landing pages rather than of the documents, many of which remain reachable by direct address. The removal was reported contemporaneously by Bloomberg under the headline 'Pentagon Removes Decades of Weapon Test Reports From Public View' (5 August 2026). These are press reports of an administrative act; the paper records the reporting and has not obtained any departmental statement of the decision itself.

¹³⁰*ibid*, giving the Department's reported rationale that adversaries could exploit artificial intelligence "and other tools to compile disparate, unclassified data points" from the reports "into highly detailed vulnerability profiles". The wording is as reported by that source and the paper has not located the departmental document it paraphrases.

¹³¹Letter from Senator Elizabeth Warren and Representative Donald Norcross to the Secretary of Defense, 12 August 2026 <https://www.warren.senate.gov/wp-content/uploads/2026/08/Final-Letter-on-GAO-review-of-DOTE.pdf> accessed 26 August 2026, citing US Government Accountability Office, *Weapon Systems Testing: Reorganization of DOD's Office of the Director, Operational Test and Evaluation*, GAO-26-108859 (30 June 2026) <https://www.gao.gov/products/gao-26-108859> accessed 26 August 2026. The quoted finding on undocumented shortfalls and the figure of 173 programmes on the 2025 oversight list are the audit's, as the letter reproduces them; the percentages for staffing and budget are the letter's characterisation. The letter's figure for authorised civilian manpower and the figure given in the audit's own summary did not agree on the readings consulted, and neither is relied on here.

¹³²Letter from Senator Elizabeth Warren to the Secretary of Defense, 1 February 2022 [https://www.warren.senate.gov/imo/media/doc/2022.02.01%20Letter%20to%20DoD%20re%20weapons%20report%20classification%20\(1\).pdf](https://www.warren.senate.gov/imo/media/doc/2022.02.01%20Letter%20to%20DoD%20re%20weapons%20report%20classification%20(1).pdf) accessed 26 August 2026, recording an announcement of December 2021 that only a controlled-unclassified version of the annual report would be made available, and stating that twenty-two programmes had information redacted and one was removed from the report entirely.

¹³³See Part IX §49 and the report cited there.

what the aggregate shows. The paper takes no position on the decision and impugns no motive. It records only that the proposition underlying the stated reason—that the open record is combinable, and that its combination is more informative than its parts—is one the Department and this paper appear to hold in common, and that it is an argument for drafting rules against what is calculable rather than against what happens to be visible in any given year, because visibility is administrable and calculability is not.

39.3 The asymmetry of error, which is the only recommendation this section makes

A rule-maker drafting under a record known to be incomplete, and now known to be *variably* incomplete, faces a question that is not about intelligence and is not answerable by better intelligence. It is a question about which way to err.

The two errors are not symmetrical, and the asymmetry runs one way.

A rule that reaches less than what exists is expensive to correct. Correcting it requires reopening an instrument, which requires a mandate, consensus among the parties, and years—§51 records that the present effort has run since 2024 under a mandate that has not yet produced a recommendation to negotiate. In the interval the conduct the rule failed to reach is lawful, and the practice that develops in that interval becomes the baseline against which any later constraint is argued.

A rule that reaches more than what exists is cheap to carry. A constraint on a function nobody has built binds nobody, costs nothing, and requires no correction; if the function is later built, the constraint is already in place, which is the only circumstance in which a rule of this kind arrives in time. The cost of over-reach is borne only if the constraint also catches something legitimate, which is a drafting question about the scope of the function rather than an argument for narrowing the reach.

This paper's own subject supplies the worked case, and it is why the asymmetry is stated here rather than left as a generality. §21.2 prices the mechanisms in question, and the pricing is what makes the comparison tractable: on the figures there, requiring the mechanism that works when nothing reaches the system costs approximately nothing, whether or not any party has built the object it would constrain. A requirement that costs approximately nothing to meet, and that binds nothing until the thing exists, is the cheapest possible instance of erring wide. §51 records that it is the requirement the current text no longer names.

The recommendation is therefore narrow and it is the only one this Part makes. **Where the record is known to be incomplete and known to move, and where the cost of a constraint can be computed and is small, the reasonable direction of error is to draft for the capability the physics permits rather than for the capability the current year's documents happen to describe.** §35 and §36 establish that the physics is open, decades old, and calculable by anyone who can read the published relations. It is the one part of the evidentiary base that no administrative decision can withdraw, and that is its whole value to a rule-maker.

40. What the compilation looked for and could not find

This section records what was searched for and not located, grouped by the kind of document or search involved. Each entry is a fact about a search or a document; none is a claim about what exists, and none speculates about why anything was not found. A rule-maker is better served by a

precise statement of what is absent from the public record than by silence about it, because a rule drafted against what is visible fails exactly where nothing is visible.

Programme and budget documents. The Army narrative justification volumes were consulted. Programme element 0604531A appears in the FY2027 volume for Budget Activity 04 in the list of programme terminations, its projects CQ5 and CQ6 recorded there as transferring to procurement and sustainment; that is a statement about a list in a published volume and nothing more. Budget justification entries for three named research programmes were not located in the volumes searched.

Technical literature. Four conference or journal papers from one major optics publisher, whose platform interposes an automated-access control, were supplied to the author from his own copies and read in full; the control was not circumvented. The propositions they carry at §15.1, §15.5.1 and §36 are now quoted from their texts rather than from their bibliographic records, and one figure previously attributed to unnamed related work has been withdrawn as a result. Two further papers on foliage penetration remain unobtained and nothing rests on them. One table in a 2014 National Research Council report could not be located.

Patents. A search of one patent classification returned no results for a particular sensing method combined with that classification; a related search returned thirteen results, all of which concerned a different application. No granted patent at any major prime contractor was found placing that sensing method on a ground weapon station. A search against one contractor holding an \$11.95m research contract returned zero patent results. Year-by-year filing counts could not be obtained; this reflects the search methods available and not the record.

Physical parameters. No numeric extinction coefficient at one specific wavelength could be extracted from the obscurant literature located.

Programme status. The sensing method used by one named ballistic-correction programme is not stated in any accessible source; across every source located, the subsystem is described only functionally. One programme's published page no longer resolves. Where secondary literature asserts that a status "is classified" without giving a source, that is recorded here as a fact about the secondary literature.

Curvature, Coriolis and spin-drift, stated explicitly because a reader will ask. No source located in this compilation names earth-curvature, Coriolis or spin-drift correction as a module in any fielded or funded system. The research to support such a claim does not exist in the material compiled here, and the paper makes none.

For a rule-maker the list carries one lesson. A constraint drafted against named programmes or named technologies will not hold, because the naming is unstable and the record is uneven: programmes are renamed, entries drop from one budget cycle to the next, and pages that once resolved now return errors. A constraint drafted against function and against use does not depend on any of this being visible, and holds at precisely the points where the naming does not.

41. What the public record establishes, stated plainly

The sections above are written so that every sentence has a document as its subject, and the discipline is the point. What the documents establish can nevertheless be stated in ordinary words, and for the audience the companion papers address it should be.

Everything the specified object requires already exists. The functions the architecture at Part V combines—an uncrewed platform, a remote weapon station, the fire control that aims it,

target recognition, photon-counting sensing—are separately named, funded and contracted in the published budget and award record, with entries current to 2026 (§32); ballistic computation is a catalogue feature and a funded project line (§33); and the physics is decades-old open literature and a NATO standard (§35, §36).

Most of it is cheap. A general-purpose legged platform is listed to the public at about thirteen and a half thousand dollars, an armed commercial quadruped at under ten thousand, and the software substrate is open source (§34). The part that is not cheap is the discriminating tier, which §21 prices at 91 per cent of the energy budget and which holds its place in effort and in price (§21.3).

The components are offered for civilian sale. That is a narrower proposition than lawfulness and it is the one the record supports: an armed commercial robot is offered for sale with a vendor’s own assertion that it is lawful in forty-eight states, and a platform with a public software development kit is offered to anyone who orders it (§34). The paper has found nothing contradicting the vendors’ assertions and has not verified them; state, municipal, export-control and end-use restrictions apply to such purchases and this paper has not surveyed them. **What the record establishes is availability at a listed price, not permission,** and a rule-maker should use it for the first and not the second.

And the safeguard is a removable layer. In the fielded stations whose reporting the paper has read, the sensing, the tracking and the firing solution are the machine’s, and the person’s approval is placed on top of them as a mode; the vendors of the armed commercial platforms say the same of their trigger (§15.7, §34). A safeguard that is a setting can be unset. The one specified at Part V is not a setting, because it is architecturally independent of the sensing it restrains, and that difference is what the companion papers need from the engineering: a safeguard that exists as a property of the machine rather than as a mode a vendor or an operator selects.

And one proposition about controls follows from the record rather than from the design. *A control tied to a named sensor or model is not function-bound. If another modality can supply the same cue, classification, selection or pose function, the function survives the component change.* The record in this Part is a record of *functions* being separately funded, contracted and sold—ranging, recognition, fire control, weapon stations—and a function that several components can supply is one that no component-specific control reaches. **Controls intended to survive substitution must therefore attach downstream, to the common function and the authority gate of §15.5.4, and not to the part.** This is the same conclusion §49 reaches from the development sequence and §15.5.3 reaches from the design, and its appearance here is what makes it a property of the record rather than an argument about it.

What the record does not establish, the paper does not say: which arrangement any party has built, how well any of it works, or what any purchaser has done with what is on sale. That boundary is the record’s own (§38, §40), and it is why the paper’s contribution to the companions is a record rather than an argument. Both published companion papers state that the object they analyse is not hypothetical in its components. Part VIII is the evidence for that statement, at the level at which evidence can be given.

Part IX—The Function Chain

42. What this Part does, and the unit it uses

Part VIII compiled what is funded, what is on sale and what the physics permits. It says nothing about what has been *built and used*, because the budget and award records do not carry that and the compilation rule forbade inferring it. This Part supplies the missing class of evidence from a different record: systems that entered service, whose functions their own governments and manufacturers have described, and whose dates are documented.

The unit of analysis here is the **function**, not the system. That choice is the whole method and it is worth defending before anything rests on it.

A system is a bundle. It carries a name, a nationality, a procurement history and a political valence, and every one of those attracts argument that has nothing to do with what the object does. A function is narrower and more durable: *waiting unattended for a month, matching an observed signature against a stored profile, engaging without a person designating the particular object*. Functions can be dated, sourced and compared across systems that have nothing else in common, and—this is the point for a rule-maker—a constraint written against a function does not go stale when the naming changes, which §40 established that it does.

The approach has a precedent worth naming, because a method presented as new should say where it has been used before. A 1998 thesis at a United States military graduate school decomposed the same problem into a set of criteria a sensor system must satisfy to support what it called precision engagement—delivery, discrimination, secrecy, dissemination, fusion, interrogation and endurance—and assessed three campaigns against them.¹³⁴ Three of those criteria are axes this paper also uses, under other names: discrimination is what Part VI prices, secrecy is the emission posture of Part IV, and endurance is the emplaced-day of §16.3. What is offered as new here is therefore not the decomposition but what is done with it: dating each function to a fielded programme, and pricing the functions against one another in the unit of §16.3.

So each section below names a function, states when that function was first fielded, and cites the systems that evidence it. The systems are the evidence; the function is the finding. No section proposes that any of these functions be combined with any other, and the paper draws no conclusion of that kind: the sections are separately sourced and are to be read separately, in the same way that §32's budget lines were.

Three limits bound the whole Part, and they are the same limits Part VIII adopted.

The sources are official descriptions and manufacturers' descriptions of their own products. Where a government test report, a service fact file or a national forces manual describes a function, that description is quoted and attributed. Where a manufacturer describes its own product, the description is marked as the manufacturer's. Neither is an independent measurement, and where the two conflict—as they do at §44 and §49—the conflict is shown rather than resolved.

Nothing here establishes performance. That a system is described as classifying targets does not establish that it classified them well, and several entries below record programmes that were terminated, that failed their own requirements, or whose capability was withdrawn. Performance

¹³⁴Haider (n 92), at the figure headed 'Sensor Criteria for Precision Engagement'. The thesis is cited here for its decomposition and for the false-activation record at §43; this paper reached its own decomposition independently and records the precedent on discovering it.

is not the question this Part asks. The question is whether the function has been *attempted at fielding scale*, because that is what bears on whether a rule needs to reach it.

Dates are first-fielded or first-documented dates, not dates of invention. They are given because the age of a function is the single most useful fact about it for the present purpose, and because the discussion this paper’s companions address is conducted as though these functions were prospective.

43. Waiting: unattended persistence with a live effect

The first function is the one this paper’s Part VI is about, and it is the oldest.

An emplaced weapon that waits is not a new object. What distinguishes the systems below from a mine is that each carries a sensing system that remains alive across the wait, drawing power continuously, and each was specified for a wait measured in weeks or months rather than in the seconds a fuze is live.

The **Mk 60 CAPTOR**, a United States encapsulated torpedo mine which entered service in 1979, was moored at depth and lay dormant until its sensing system detected a target; published descriptions state that it could remain on station for a period described as weeks to months, and it was withdrawn from the inventory in the years after 2001.¹³⁵ The **TM-83**, a Soviet off-route anti-armour mine of 1983, is described in the reference literature as holding a seismic channel in a standby mode and waking an infrared channel on a detection, with an action time of thirty days before a battery change is required.¹³⁶ The **M93 Hornet**, a United States wide-area munition delivered from 1999, was required to hold an active life of at least thirty days with a desired figure of 180, and a government audit records that the programme did not achieve that requirement.¹³⁷ The **PTKM-1R**, a Russian top-attack anti-vehicle mine documented in Ukraine from April 2022, is specified by its exporter with a settable self-destruction time of one to ten days.¹³⁸

Two things follow and the second is the one that matters here.

¹³⁵On the Mk 60 CAPTOR encapsulated torpedo mine see GlobalSecurity, ‘MK 60 Encapsulated Torpedo (CAPTOR)’ <https://www.globalsecurity.org/military/systems/munitions/mk60.htm> accessed 26 August 2026, stating that the weapon “lies dormant until a target is detected”; and National Research Council, *Naval Mine Warfare: Operational and Technical Challenges for Naval Forces* (National Academies Press 2001) ch 4, recording the weapon as a moored mine at medium depth and the majority of the inventory as being withdrawn. Service entry is given as 1979 in the reference literature; one secondary account gives 1976, and the difference is not resolved here. The endurance figure is reported as weeks to months and no exact figure could be sourced.

¹³⁶On the TM-83 see the reference description at https://en.wikipedia.org/wiki/TM-83_mine accessed 26 August 2026, giving a two-stage seismic-standby and infrared-active arrangement and an action time of thirty days to a battery change. This entry rests on a single reference source of the weakest class used anywhere in this paper; it is cited for the existence and shape of the arrangement and no figure from it is relied on elsewhere.

¹³⁷US Department of Defense, Office of the Inspector General, *Acquisition of the Wide Area Munition*, Report No D-2002-011 (24 October 2001) <https://media.defense.gov/2001/Oct/24/2001714099/-1/-1/02-011.pdf> accessed 26 August 2026, recording the requirement that the munition self-destruct “just prior to battery life expenditure” with an active life of at least thirty days and a desired 180 days, and recording that the test authority concluded the requirement was not achieved.

¹³⁸Rosoboronexport, ‘PTKM-1R’ product page <https://roe.ru/en/production/land-forces/engineer-equipment/ptkm-1r/> accessed 26 August 2026, giving a self-destruction time of one to ten days; and N R Jenzen-Jones and others, ‘Russian PTKM-1R top-attack anti-vehicle mine documented in Ukraine’ (*Armament Research Services*, 2022) <https://armamentresearch.com/russian-ptkm-1r-top-attack-anti-vehicle-mine-documented-in-ukraine-2022/> accessed 26 August 2026, recording documentation of the mine in Ukraine from April 2022 and giving the same self-destruction range. The exporter’s figures are a manufacturer’s published specification.

The wait is bounded by energy in every case, which is what Part VI predicted and is the reason that Part exists: the TM-83's thirty days is a battery-change interval, the Hornet's requirement was written as self-destruction "just prior to battery life expenditure", and the PTKM-1R's ten days is an outer figure on a device whose sensing is live throughout. **The endurance of an emplaced sensing weapon has always been an energy specification, and it has been written as one by the organisations that fielded these systems.** What §16.3 proposes is not a new kind of quantity. It is the quantity these programmes were already using, stated in a unit that permits architectures to be compared rather than left implicit in a battery-change interval.

And the durations are in the same range this paper's design case occupies. A thirty-day armed emplacement is not an extrapolation. It is the figure two of these four programmes were specified against, and it is the figure Amended Protocol II fixes for the self-destruction of a remotely delivered anti-personnel mine.

44. Classifying: matching an observed signature against a stored profile

The second function is the one §21 prices, and the record shows it fielded three decades before this paper.

The **M93 Hornet** used acoustic and seismic sensors in its ground platform "to detect, track, and classify potential targets", with a passive infrared sensor on the submunition; the reference literature records that the munition disregarded any target below a stated mass threshold, prioritised among those remaining, and predicted the movement of the one it selected.¹³⁹ The **Mk 60 CAPTOR** is described as carrying a passive acoustic targeting system that "could discriminate between hostile submarines and surface vessels and friendly submarines that might also be operating in the area"—which is to say that it distinguished not merely target class but target *allegiance*, on signature, unattended, in 1979.¹⁴⁰ The **PTKM-1R** is described by the research literature as identifying targets by comparing a target's acoustic and seismic signature against an internal database, and by its exporter as selecting "only those equipment that meet the specified parameters in terms of noise and ground vibration".¹⁴¹

The distinction that matters, and that the discussion routinely loses, is between an **influence fuze** and a **classifier**. The mines of the FASCAM family—the scatterable anti-armour and anti-

¹³⁹Federation of American Scientists, 'M93 HORNET' <https://man.fas.org/dod-101/sys/land/docs/m93.htm> accessed 26 August 2026, stating that the munition "utilizes acoustic and seismic sensors in its ground platform to detect, track, and classify potential targets" and that the submunition uses a passive infrared sensor; and Forecast International, 'M93 Hornet Wide Area Munition' (archived programme report) https://www.forecastinternational.com/archive/disp_pdf.cfm?DACH_RECNO=333 accessed 26 August 2026, recording the mass threshold, the prioritisation among candidates and the movement prediction, and recording the programme's termination in 2002 after 441 units.

¹⁴⁰Joseph Trevithick, 'The U.S. Is Getting Back Into Naval Mine Warfare' (*The War Zone*, 2019) <https://www.twz.com/25235/the-u-s-is-getting-back-into-naval-mine-warfare-with-new-sub-launched-and-air-dropped-types> accessed 26 August 2026, describing a reliable-acoustic-path passive sonar targeting system that "could discriminate between hostile submarines and surface vessels and friendly submarines that might also be operating in the area"; and GlobalSecurity (n 135), describing the detection system as designed to seek hostile submarines while ignoring surface craft and friendly submarine signatures. Both are secondary descriptions of a withdrawn system.

¹⁴¹Jenzen-Jones and others (n 138), stating that the mine identifies targets by comparing a target's acoustic and seismic signature with an internal database; Rosoboronexport (n 138), stating that the system selects only equipment meeting specified parameters of noise and ground vibration. A secondary research account and a manufacturer's page respectively. A humanitarian organisation's account of the same weapon gives a materially larger detection range and a different sensor complement, and the difference is not resolved here: Human Rights Watch, 'Background Briefing on Landmine Use in Ukraine' (15 June 2022) <https://www.hrw.org/news/2022/06/15/background-briefing-landmine-use-ukraine> accessed 26 August 2026.

personnel families of the same era—used magnetic influence fuzing, which reacts to a change in a field, and trip wires, which react to contact.¹⁴² A magnetic fuze performs no comparison and holds no profile: it is a threshold on a single channel. A system that holds a stored signature library, compares an observation against it, applies a mass threshold and ranks candidates is doing something categorically different, and it is doing the thing the law calls discrimination, at the level of target class.

This qualifies a claim made at §21.1 and the qualification strengthens it rather than weakening it. Discrimination in the sense of *distinguishing one class of materiel from another by signature* has been fielded since the 1970s and was never unavailable. What has never been fielded in an emplaced weapon, and what §21.2 prices at about 1.1 Wh per emplaced-day, is discrimination in the sense the law of targeting requires: distinguishing a combatant from a civilian, which is a judgement about a person and not a signature match against a vehicle. The mine law’s silence on discrimination is explained not by the impossibility of all discrimination but by the impossibility of *that* discrimination, and the price this paper computes is the price of the kind that was missing.

45. Selecting: engaging without a person designating the object

The third function is the one the treaty discussion is actually about, and it has been in service since the 1990s.

The distinction here is not between a machine that shoots and a person who shoots. It is between a person who designates *the particular object to be attacked* and a person who designates only a *category and an area* and delegates the choice of object to the machine. The second is what the text analysed at §51 calls the identification, selection and engagement functions.

The **IAI Harpy**, an anti-radiation loitering munition in service by 1994 on the mapping of the Stockholm International Peace Research Institute, searches an area for emitters matching a stored profile and attacks what it finds. The person launching it directs it to a region believed to contain such emitters; they do not direct it against any particular one.¹⁴³ SIPRI’s survey identified only four operational systems capable of finding, tracking and attacking targets in complete autonomy once launched, and recorded that the Harpy—the oldest of them—“operates in complete autonomy”, while its successors “include both a human-in-the-loop and fully autonomous mode”.¹⁴⁴

That last clause carries more weight than anything else in this Part, and §49 returns to it. The direction of travel in this family was not from human control toward autonomy. It was from autonomy toward an added human-control mode.

¹⁴²US Army, *FM 20-32: Mine/Countermining Operations*, ch 3, on the scatterable mine families, giving magnetic influence fuzing for the anti-armour types and trip-wire actuation for the anti-personnel types, and self-destruct settings of four hours, forty-eight hours and fifteen days according to type <https://www.globalsecurity.org/military/library/policy/army/fm/20-32/chap3.html> accessed 26 August 2026.

¹⁴³Vincent Boulanin and Maaike Verbruggen, *Mapping the Development of Autonomy in Weapon Systems* (Stockholm International Peace Research Institute 2017) https://www.sipri.org/sites/default/files/2017-11/sipri_report_mapping_the_development_of_autonomy_in_weapon_systems_1117_1.pdf accessed 26 August 2026, recording the Harpy in the dataset with a service date of 1994 or earlier and describing it as capable of finding and attacking radar installations autonomously. The characterisation of what the launching person does and does not designate follows Paul Scharre, ‘Autonomy, “Killer Robots,” and Human Control in the Use of Force’ (*Just Security*, 9 July 2014) <https://www.justsecurity.org/12708/autonomy-killer-robots-human-control-force-part/> accessed 26 August 2026.

¹⁴⁴Boulanin and Verbruggen (n 143), recording four operational systems able to find, track and attack in complete autonomy once launched, that the Harpy is the oldest and operates in complete autonomy, and that the later systems include both a human-in-the-loop and a fully autonomous mode.

The **M93 Hornet** performed the same function on the ground and in the same period. Its own programme literature and the government’s test reporting both describe it as “fully autonomous from final arming to target engagement”.¹⁴⁵ A person emplaced it and set its arming; no person chose the vehicle it struck.

46. Accepting: autonomous engagement already fielded, already accommodated in policy

The fourth function is not a capability but a permission, and it is the part of the record most often missing from the discussion.

Autonomous engagement is not a prospect awaiting a first deployment. It is a fielded practice of long standing, and one state has written down an explicit policy accommodation for a defined subset of it. What follows is a claim about practice and about policy. It is **not** a claim that autonomous engagement has been determined to be lawful: a national directive is policy rather than law, exempting a class of system from a *senior review* requirement is not exempting it from legal review, and the directive itself preserves the weapons review, testing and law-of-war obligations that apply to everything else. What long practice without recorded objection supplies is evidence bearing on lawfulness, not a determination of it, and this Part does not treat the two as the same thing.

The **Phalanx** close-in weapon system entered service with the United States Navy in 1980. The Navy’s own fact file describes it as the only deployed close-in weapon system “capable of autonomously performing its own search, detect, evaluation, track, engage and kill assessment functions”, and states that it “automatically detects, evaluates, tracks, engages and performs kill assessment” against the threats it is built for.¹⁴⁶ The **Aegis** combat system, first deployed in 1983, is described by the same service as designed “as a total weapon system, from detection to kill”, operated through doctrine statements which “define automatic actions for targets meeting specific conditions”.¹⁴⁷ The **Counter-Rocket, Artillery and Mortar** system has been described by the United States in its own statement to this treaty body as an example of a system that would “select and engage incoming projectiles automatically”.¹⁴⁸

¹⁴⁵Federation of American Scientists (n 139); and Director, Operational Test and Evaluation, FY2001 annual report, Wide Area Munition, reproduced at <https://www.globalsecurity.org/military/library/budget/fy2001/dote/army/01wam.html> accessed 26 August 2026. Both describe the munition as “fully autonomous from final arming to target engagement”. A programme report of the same system separately describes it as remotely operated by a control unit; the control unit provided arming, deactivation and command destruction and did not place a person in the engagement decision, and the two descriptions are consistent on that reading.

¹⁴⁶US Navy, ‘MK 15 Phalanx Close-In Weapon System (CIWS)’ fact file <https://www.navy.mil/resources/fact-files/display-factfiles/article/2167831/mk-15-phalanx-close-in-weapon-system-ciws/> accessed 26 August 2026. Service entry is given as 1980.

¹⁴⁷US Navy, ‘Aegis Weapon System’ fact file <https://www.navy.mil/Resources/Fact-Files/Display-FactFiles/Article/2166739/aegis-weapon-system/> accessed 26 August 2026; and on doctrine statements defining automatic actions for targets meeting specified conditions, GlobalSecurity, ‘AEGIS Core Processes’ <https://www.globalsecurity.org/military/systems/ship/systems/aegis-core.htm> accessed 26 August 2026.

¹⁴⁸United States statement at the Group of Governmental Experts on Lethal Autonomous Weapons Systems, agenda item 5(d), 5 August 2021 <https://geneva.usmission.gov/2021/08/05/u-s-statement-at-the-gge-on-laws-during-the-discussion-of-agenda-item-5d/> accessed 26 August 2026, giving the Counter-Rocket, Artillery, Mortar system as an example of a system that would “select and engage incoming projectiles automatically”. See also the United States statement of 27 March 2019 to the same body, naming the Phalanx, Aegis and Patriot systems as having autonomous functions <https://geneva.usmission.gov/2019/03/27/convention-on-certain-conventional-weapons-consideration-of-the-human-element-in-the-use-of-lethal-force/> accessed 26 August 2026.

The scale is not marginal. A survey by the Center for a New American Security recorded that at least thirty nations operate human-supervised defensive systems of this class.¹⁴⁹

And the policy accommodation is explicit, and narrow. The United States Department of Defense directive on autonomy in weapon systems, in its 2023 revision, exempts from its senior review requirement those operator-supervised autonomous weapon systems used “to select and engage materiel targets for local defense to intercept attempted time-critical or saturation attacks”, including for “static defense of installations with personnel, including networked defense where the autonomous weapon system is not co-located with the installation”.¹⁵⁰

Two observations follow, and a paper addressed to this audience should make both plainly.

The line already drawn in practice is a line about the target, not about the autonomy. The directive’s exemption is for *materiel* targets; the International Committee of the Red Cross recommends prohibiting autonomous weapon systems “designed or used to apply force against persons” and constraining the remainder to “objects that are military objectives by nature”.¹⁵¹ A defence ministry and a humanitarian institution, arguing from opposite premises, converge on the same boundary. What is contested is not whether a machine may select and engage. It is what it may select and engage.

And the architecture this paper specifies sits precisely on the exempted side of that boundary as the directive draws it—on the boundary’s *geometry*, that is, and not on its target class—a static emplacement, defending ground, networked, not co-located with its operator—while the target it is specified against is a person walking. That is the single most consequential sentence in this Part for a delegate reading it, and the paper states it as a fact about where the existing accommodation falls rather than as an argument about where it should.

47. Re-tasking: changing an emplaced weapon’s state from the rear

The fifth function is the one §30 is about, and the record required this paper to withdraw a claim of priority.

The **M131 MOPMS** is a United States munition dispenser whose hand-held radio control unit permits one person to control as many as fifteen groups of modules from a remote location at a range given as 300 to 1,000 metres, to command their destruction, and to recycle a four-hour self-destruct time three times.¹⁵² The **M7 Spider**, fielded from around 2009, is a hand-emplaced

¹⁴⁹Paul Scharre and Michael C Horowitz, *An Introduction to Autonomy in Weapon Systems* (Center for a New American Security, February 2015) 12, recording that at least thirty nations use human-supervised defensive systems with greater autonomy, with the systems and states listed at the appendix.

¹⁵⁰US Department of Defense Directive 3000.09, *Autonomy in Weapon Systems* (25 January 2023) para 1.2.d(2) <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf> accessed 26 August 2026. The 2023 revision also defines an autonomous weapon system as one which, once activated, can select and engage targets without further intervention by an operator, and states that this includes operator-supervised systems designed to allow operators to override operation but which can select and engage without further operator input after activation: *ibid*, Glossary G.2. The paper cites the directive for what it provides and takes no position on it.

¹⁵¹International Committee of the Red Cross, *ICRC Position on Autonomous Weapon Systems* (12 May 2021) https://www.icrc.org/sites/default/files/document_new/file_list/icrc_position_on_aws_and_background_paper.pdf accessed 26 August 2026, recommending prohibition of autonomous weapon systems designed or used to apply force against persons, and limits on the types of target constraining them to objects that are military objectives by nature.

¹⁵²Federation of American Scientists, ‘M131 Modular Pack Mine System (MOPMS)’ <https://man.fas.org/dod-101/sys/land/m131.htm> accessed 26 August 2026, giving control of as many as fifteen groups of modules at 300 to 1,000 metres, a factory-set four-hour self-destruct time recyclable three times, and command destruction. The

networked munition under remote command, whose components including undetonated munitions are recoverable and reusable.¹⁵³ Its programme successor, the United States Army’s terrain-shaping obstacle effort, is described in the Army’s own publications as carrying “on-off-on, remote destruct, self-locate and self-reporting capabilities” with networked two-way command and control, for the stated purpose of recovering the effects of landmines while keeping them under friendly control after emplacement.¹⁵⁴

An emplaced obstacle whose arming state can be changed from the rear without anyone going forward is therefore not new, and §30 no longer claims that it is. What §30 claims, and what this record does not anticipate, is set out there: these systems toggle a munition of fixed lethal effect between armed and safe, whereas the object specified at Part V has a discriminating sensor whose sector, release criteria and readiness are parameters. The mutable quantity in the prior art is arming. The mutable quantity in the specification is judgement.

48. Supervising: what a person monitoring an unattended sensor field achieves

The five functions above are machine functions. This one is the human function that the others are held to require, and it is the only function in this Part whose performance envelope has been *measured and published* rather than described.

The requirement now before the treaty body, at paragraph 35(c), is for “measures to intervene in the operation of LAWS, to the extent possible or feasible, including to deactivate or neutralize LAWS in a timely manner”, and §51 reads that against what it costs. But the requirement presupposes something upstream of cost: that a person supervising the system is in a position to know that deactivation is warranted. That presupposition has been tested.

In September 1978 the United States Army Research Institute for the Behavioral and Social Sciences published a study of the effect of workload on operators monitoring unattended ground sensors.¹⁵⁵ Experienced operators monitored activations from sensor grids under six workload conditions—two levels of target activity crossed with three grid sizes of 27, 54 and 108 sensors—and their detection completeness, false-alarm rate and estimation accuracy were recorded against ground truth. The findings are stated in the report’s own abstract:

corresponding field manual gives the total recycled duration as approximately thirteen hours where this source gives sixteen; both are recorded and neither is relied on.

¹⁵³US Army, ‘Along came a spider’ (*army.mil*, 9 October 2012) https://www.army.mil/article/90105/along_came_a_spider/ accessed 26 August 2026, recording remote control of the munition field, that target identification must be realised prior to detonation, and that all components including undetonated munitions are recoverable and reusable.

¹⁵⁴US Army, ‘Critical events in developing the concept of operations for terrain-shaping obstacles’ (*army.mil*, 2019) https://www.army.mil/article/224834/critical_events_in_developing_the_concept_of_operations_for_terrain_shaping_obstacles accessed 26 August 2026, recording networked two-way command and control and “on-off-on, remote destruct, self-locate and self-reporting capabilities”; and US Army, ‘Soldiers shape the battlefield with next generation of terrain-shaping munitions’ (*army.mil*, 2019) <https://www.army.mil/article/225611> accessed 26 August 2026.

¹⁵⁵Lawrence R Edwards, Sterling S Pilette and Billy E Biggs (HRB-Singer, Inc.) and Harold Martinek, *The Effect of Workload on Performance of Operators Monitoring Unattended Ground Sensors*, Technical Paper 321, US Army Research Institute for the Behavioral and Social Sciences, Battlefield Information Systems Technical Area (September 1978), DTIC AD-A061694, read in full. Target activity level was defined as either 5 or 14 vehicle groups per 27 sensors per hour, and operators monitored 27, 54 and 108 sensors in counterbalanced sequence. The quoted passages are from the report’s abstract. The study concerns operators monitoring reporting sensors and does not concern authorisation of force.

“Percentage of detections decreased linearly from 85% (4 to 6 vehicle groups per hour presented on 27 sensors) to 26% (about 40 vehicle groups per hour presented on 108 sensors).”

“Without special training, operators should monitor no more than 60 sensors being activated by a maximum of 10 vehicular groups per hour (detection completeness for this condition was 76%).”

Two further results are recorded and the second is the more useful. Operator accuracy in estimating a target’s direction degraded significantly with workload. But **the number of false alarms did not significantly change over workload**—which is to say that the operator’s error under load was not that they saw things which were not there, but that they *failed to see things which were*.

A companion study by the same team, published a month earlier, tested whether training closes the gap.¹⁵⁶ It found that an individualised training programme significantly improved detection, identification, speed and quantity estimation, and that a job aid improved most of the same measures again. It also recorded what trained operators achieved at the most favourable load tested—strings of three and four sensors:

“Operator performance on 3- and 4-sensor strings was 77% detection completeness with virtually no false alarms.”

The two studies are not directly comparable, being different tasks and scenarios, and neither is offered as a measurement of the other. What survives their difference is the shape they share. Training helps and is worth doing; the residual omission remains of the order of a quarter even for trained, aided operators at the lightest load tested; and in both studies the error is omission rather than false alarm.

Three things follow, and the paper states them as facts about a study rather than as an argument.

Training improves the function and does not remove its ceiling. That is the finding of the two studies read together, and it matters for how a supervisory requirement should be read: a rule that assumes shortfalls in human oversight are a training problem is assuming something the army that fielded these systems tested and did not find.

The human supervisory function has a measured capacity, and it is small. Fifty-nine per cent of detection completeness is lost between the lightest and heaviest conditions tested, and the study’s own recommendation—no more than sixty sensors at ten target groups an hour—is a bound on a sensor field far smaller than the ones §32 records being funded and §34 records being sold. A supervisory requirement written without a stated ratio of sensors to operators is a requirement whose satisfaction cannot be assessed.

The degradation is in omission rather than commission, which is the direction that matters here. A supervisor who generates false alarms under load is visibly failing. A supervisor who misses three-quarters of what crosses is failing invisibly, and the record they generate looks like a quiet sector. §23’s third risk—that a closed system generates no contemporaneous record

¹⁵⁶Sterling Pilette and Bill Biggs (HRB-Singer, Inc.) and Harold Martinek, *The Value of Special Training and Job Aids for Improving Unattended Ground Sensor Operator Performance*, Technical Paper 304, US Army Research Institute for the Behavioral and Social Sciences, Battlefield Information Systems Technical Area (August 1978), DTIC AD-A061046, read in full. The quoted passage is from the report’s abstract. The study is an interim report covering work from November 1972 to May 1975 and concerns operators monitoring reporting sensors; as with (n 155) it does not concern authorisation of force.

anyone can read—has a counterpart here: an overloaded supervised system generates a record that is present and wrong.

And the finding is nearly fifty years old. It was measured by the army that fielded the sensors, on operators experienced with them, and published in the open literature in 1978. It is not a prediction about future systems and it is not contested by anything this compilation located. Whatever weight a rule places on human supervision, the capacity of that supervision has been a measurable quantity for half a century, and this paper’s only observation is that the current text places weight on it without naming the quantity.

One boundary. The study measured monitoring of *reporting* sensors, not supervision of a system with effectors, and the operators’ task was detection and estimation rather than authorisation of force. Nothing here establishes what a person authorising an engagement would achieve, and this paper does not extrapolate from the one to the other. What it establishes is narrower and is enough: that the performance of a human supervising an unattended sensor field is a function of how much they are given to supervise, that the function has been measured, and that it declines steeply.

49. The finding: autonomy is a mode, and the safeguard is what was added

The five functions above are separately sourced and separately dated. Read together they yield one finding, and it is this Part’s contribution.

In every fielded system in this record for which the sequence is documented, the autonomous capability came first and the human control was added afterwards, as a mode.

A second sense of “mode” belongs beside it and must not be allowed to become a new taxonomy. The *classifying* function of §44 can be implemented by matching a seismic signature against a stored profile, by a camera and a trained model, by a point-cloud detector working on lidar returns, or by a fusion of several—and the function is the same in each. Modality is therefore also a mode, in exactly the sense that autonomy is: a way the same function is delivered, switchable without changing what the system is for. **The point is functional substitutability and not a new label.** A rule that reached “camera-based target recognition” would not reach the same function delivered from a point cloud, and a rule that reached “lidar-equipped systems” would not reach it delivered from a camera. §41’s proposition and §15.5’s gate are the two halves of the answer: describe the function, and place the control downstream of every path that can perform it.

Switching modality is not a human decision and does not restart the interval. An automatic transition from a passive to an active sensing path changes the evidence and may change the usable envelope, but it is not an act of authorisation, and §53 states the consequence formally.

The evidence is not inferential. It is in the manufacturers’ and the governments’ own words.

Of the **Super aEgis II**, a static armed turret which its manufacturer states detects and tracks to about three kilometres, a senior engineer of the company told a reporter in 2015: “It wasn’t initially designed this way. Our original version had an auto-firing system. But all of our customers asked for safeguards to be implemented.” He added: “Technologically it wasn’t a problem for us. But they were concerned the gun might make a mistake.”¹⁵⁷ The safeguard is, on the maker’s own

¹⁵⁷Simon Parkin, ‘Killer robots: The soldiers that never sleep’ (*BBC Future*, 16 July 2015) <https://www.bbc.com/future/article/20150715-killer-robots-the-soldiers-that-never-sleep> accessed 26 August 2026, quoting Jungsuk

account, a configuration retrofitted to an auto-firing baseline at the customer’s request, and its removal was described as commercially rather than technically constrained.

Of the **Patriot** air and missile defence system, a human-factors analysis published by an author from the United States Army’s own research lineage sets out the two modes in a table: “Semi-automatic mode: Human operator must authorize engagement or system will not fire. Automatic mode: System will fire unless human operator halts engagement.”¹⁵⁸ One system, one switch, and two entirely different legal characterisations of the same machine.

Of the **Phalanx**, the reference description records that it fires automatically at a set range “or recommends fire to the operator”—alternative behaviours of one architecture.¹⁵⁹

Of the **Harpy** family, SIPRI records that the successors added a human-in-the-loop mode which the original did not have.¹⁶⁰

And of the **M7 Spider**, the government’s own test authority reported in 2010 that “Spider no longer has the capability to engage targets autonomously. All engagements use ‘man-in-the-loop’ control to engage targets.”¹⁶¹ That sentence is the most useful in the Part, because it records the operation running in the other direction: a fielded system whose autonomous engagement capability was *removed*, as a matter of national policy, and whose removal was documented. What can be removed by policy can be required to be absent by law. The function is regulable because it is separable, and a government has already demonstrated the separation on a real programme.

The consequence for this paper’s own specification is stated at §15.7 and is repeated here in one sentence because this Part is the evidence for it. A safeguard that is a mode can be unset by whoever holds the configuration, and nothing in the artefact resists it; a safeguard that is architecturally independent of the perception it restrains cannot be unset without building a different machine. The record above is what makes that distinction load-bearing rather than fastidious: it shows an industry in which the restraint is, as a matter of documented practice, the part that is added last and held loosest.

One boundary should be marked before this is over-read. That autonomy preceded restraint in these programmes is a fact about how they were engineered and sold. It is not a claim that any operator has removed a safeguard, and this paper asserts none: no source in this compilation records an instance of a fielded safeguard being disabled in use, and §50 records that absence with the others.

Park, a senior research engineer of the manufacturer. The passage is reproduced, with attribution to the BBC, in Erico Guizzo and Evan Ackerman, ‘Do We Want Robot Warriors to Decide Who Lives or Dies?’ (*IEEE Spectrum*, 31 May 2016) <https://spectrum.ieee.org/do-we-want-robot-warriors-to-decide-who-lives-or-dies> accessed 26 August 2026, which also gives the three-kilometre figure and records the manufacturer’s statement that its systems as configured will not fire until a person confirms the target. The BBC page was not accessible when consulted, and the quotations are taken from the IEEE reproduction, which credits the original.

¹⁵⁸John K Hawley, *Patriot Wars: Automation and the Patriot Air and Missile Defense System* (Center for a New American Security, January 2017) 7. The author writes from the United States Army’s human-factors research lineage and the report is a study of two fratricides in 2003.

¹⁵⁹Reference description of the system’s engagement behaviour, recording automatic fire at a set range or a recommendation to the operator; see US Navy (n 146) for the service’s own description of the autonomous functions.

¹⁶⁰Boulain and Verbruggen (n 143).

¹⁶¹Director, Operational Test and Evaluation, FY2010 annual report, Spider XM7 Network Command Munition <https://www.dote.osd.mil/Portals/97/pub/reports/FY2010/army/2010spider.pdf> accessed 26 August 2026: “Spider no longer has the capability to engage targets autonomously. All engagements use ‘man-in-the-loop’ control to engage targets.” The change followed the United States landmine policy of 2004.

50. What this record does and does not contain

Four findings, recorded because a rule-maker is better served by a precise statement of what is present and absent than by silence. The first was a negative in earlier versions of this paper, and it has been overtaken by events in a way that bears directly on the argument.

A documented case of autonomous target selection in armed conflict now exists, and this paper's previous negative finding is withdrawn. On 6 July 2026 a Russian Molniya strike drone hit a filling station in Zaporizhzhia, killing three civilians. Ukrainian investigators recovered the wreckage of that airframe and of others of the same type, and reported that each carried an Nvidia Jetson Orin compute module running target-recognition software; the modules were not encrypted, and the terrain imagery and targeting code were readable. On the reported account the operators programmed a destination and a general target area before launch, and the onboard system selected the particular object at the site—a propane tank, from a class it had been trained to recognise. The case was published on 24 August 2026 and the finding is attributed to Ukrainian drone specialists, military officers and forensic investigators, with an outside analyst describing it as the first documented case of civilian deaths caused by a Russian drone using fully autonomous targeting.¹⁶²

What it establishes, stated narrowly. It establishes autonomous selection of a *material object* within an area a person designated. It does not establish autonomous selection of persons; the reported classes are containers and infrastructure, and the deaths were of people near the object rather than of people the system chose. The renewed call of 25 August 2026 by the Secretary-General and the President of the International Committee of the Red Cross is consistent with that reading.¹⁶³ **The negative this paper can still defend, and the one a delegate should use, is therefore the narrower one: no independently documented case of a weapon system autonomously selecting a *person* as its target in armed conflict was located as at 26 August 2026.** That is a materially different statement from the one this paper made in earlier versions, and the difference is exactly the boundary §46 identifies as the one on which a defence ministry's own directive and the ICRC converge.

¹⁶²The case was published by *The New York Times* on 24 August 2026 and reported widely the same and following day; see, among others, Meduza, 'NYT: A fully autonomous Russian AI drone run by an Nvidia minicomputer killed 3 civilians in the Ukrainian city of Zaporizhzhia' (24 August 2026) <https://meduza.io/en/feature/2026/08/24/nyt-a-fully-autonomous-russian-ai-drone-run-by-an-nvidia-minicomputer-killed-3-civilians-in-the-ukrainian-city-of-zaporizhzhia>; Euromaidan Press, 'A Russian drone chose its own target and killed three. The code that told it what to look for was readable' (25 August 2026) <https://euromaidanpress.com/2026/08/25/a-russian-drone-chose-its-own-target-and-killed-three-the-code-that-told-it-what-to-look-for-was-readable/>; and Tom's Hardware, 'Nvidia Jetson Orin-guided Russian AI drone killed three civilians in Ukraine, forensic teams say' (2026) <https://www.tomshardware.com/tech-industry/drones/nvidia-jetson-orin-guided-the-russian-ai-drone-that-killed-three-civilians-in-ukraine-forensic-teams-say>, all accessed 26 August 2026. The strike is reported as having occurred on 6 July 2026 at a filling station in Zaporizhzhia; the attribution of autonomous target selection is to Ukrainian investigators examining recovered airframes, with the characterisation as a first documented case attributed to Kateryna Bondar of the Center for Strategic and International Studies. This paper records the case as reported and has not examined the wreckage, the code or the investigators' methods; the reporting is recent, single-origin in its primary investigation, and should be treated as such. Nothing in this paper's argument depends on the case being correct in every particular: if it is withdrawn, the negative finding it displaced returns, and the evidentiary point at §50 stands on its own reasoning.

¹⁶³Renewed call by the UN Secretary-General and ICRC President to adopt rules on autonomous weapons' (International Committee of the Red Cross, 25 August 2026) <https://www.icrc.org/en/statement/renewed-call-un-secretary-general-and-icrc-president-adopt-rules-autonomous-weapons> accessed 26 August 2026. The two had previously called, in October 2023, for negotiations to be concluded by 2026.

And the case falsifies this paper’s evidentiary claim, which is the correction that matters most here. Earlier versions asserted that an autonomous and a commanded engagement “are not distinguishable from the wreckage”, and concluded that a rule waiting for a documented case would be waiting on an observation the physical evidence could not supply. That was wrong, and it was wrong in an instructive direction. **Autonomy is not visible in the effect. It is visible in the artefact.** The crater, the fragments and the casualty pattern of an autonomous engagement are indistinguishable from those of a commanded one, and that part of the claim stands. What the earlier version failed to see is that the decision does not live in the effect; it lives in a compute module, a model and a set of class labels, and those survive, and in this case they were readable because nobody had encrypted them.

Three consequences follow for a rule-maker and they are practical rather than rhetorical.

The evidentiary question is an exploitation question. A rule drafted against autonomous selection is verifiable to the extent that recovered compute hardware can be exploited, and not otherwise. That makes the recoverability and readability of the decision path—not the forensics of the effect—the thing an inspection or accountability regime would have to reach. It also means the evidence is contingent on an adversary’s engineering hygiene: this case was legible because the modules were unencrypted, and the same case with signed and encrypted firmware would have yielded a strike indistinguishable from any other.

A verification regime therefore has a design consequence it can state. If autonomy is to be provable after the fact, the decision path must be recorded in a form that survives and can be read—a log of what was selected, by which sensing path, under which configuration, and on whose authority, held in the platform and retrievable from the wreck. **A minimal logical record is a tuple per decision:**

$$\ell_k = (t_k, m_k, \theta_k, \text{id}(A_H), V_k, \Gamma_k, h(y_k), d_k)$$

—the time, the modality in use, its configuration, an identifier for the authority in force, the survey-integrity and authorisation predicates as evaluated, a protected digest or reference $h(y_k)$ for the sensor-derived evidence, and the decision d_k . **The modality and configuration fields are what make the record answer the question §15.5 raises:** a log that records only “engaged” cannot distinguish a camera decision from a point-cloud decision taken under a different threshold, and those are different systems on the materiality test at §15.5.5.

This is a logical schema and not a costing. The record is small and writing it is cheap. Everything else about it is unpriced and this paper does not price it: surviving the destruction of the platform, remaining readable afterwards, resisting alteration by the party that holds it, and binding the record to the configuration actually in force are four separate engineering problems. §29 withdraws the aggregate figure an earlier version gave for them. **The obligation to be auditable is cheap to state and its cost is unmeasured,** which is a weaker claim than the one this paragraph previously made and is the one the evidence supports.

And the compute module is the one this paper had already assumed. Part VI §17.3 takes its tier-three power class from published figures for an uncooled long-wave infrared core and an Nvidia Jetson Orin Nano edge inference module, at 5 to 25 watts under load.¹⁶⁴ The recovered airframes carried a Jetson Orin. The paper’s assumed component class is the component class actually fielded, by an adversary, in the year the paper was written—which is a check on

¹⁶⁴The component class stated at §17.3 (n 86); the reported module on the recovered airframes is of the same family (n 162).

its arithmetic that it did not expect to obtain and could not have arranged. §34's argument that the parts are ordinary, commercially available and cheap is no longer an inference from catalogue prices.

The case usually cited before this one did not carry what was placed on it, and the contrast is now the useful part. That case is a passage in the report of a United Nations panel of experts of 8 March 2021, which states at paragraph 63 that retreating forces “were subsequently hunted down and remotely engaged by the unmanned combat aerial vehicles or the lethal autonomous weapons systems” and that these “were programmed to attack targets without requiring data connectivity between the operator and the munition: in effect, a true ‘fire, forget and find’ capability”.¹⁶⁵ The annex the report cites in support of that sentence is a wreckage-identification panel: it records the first sighting of the munition in the country, states that its presence is a non-compliance with an arms embargo, and cites as its own sources a confidential source, the manufacturer's product website and a subscription reference service. It says nothing about autonomy, about an engagement, or about casualties. The panel's mandate was embargo compliance. The autonomy sentence is a characterisation of a capability, in part from the manufacturer's own marketing, and the report does not state that anyone was killed by a system operating without supervision—a reading several analysts have since made, and which the manufacturer disputes.¹⁶⁶ **The difference between that case and the Zaporizhzhia one is not the strength of the language. It is that in one of them somebody recovered the computer and read the code.** A rule-maker should take from the pair a standard of evidence rather than a count of incidents.

No source located records a fielded safeguard being disabled in use, as §49 states. The Molniya case is not such a record: nothing in the reporting indicates that a human-control mode existed on that airframe and was removed.

And the treaty body's own subsidiary instrument has not taken the question up. The compilation examined the final document of the most recent annual conference of the Amended Protocol II High Contracting Parties and found no reference to self-destruction, self-deactivation, self-neutralization, sensor-fuzed munitions, or to the withdrawals from the anti-personnel mine convention recorded at §51, which fell within its reporting period.¹⁶⁷ That is a fact about a document and not a criticism of it; its consequence is at §51.

¹⁶⁵Letter dated 8 March 2021 from the Panel of Experts on Libya established pursuant to resolution 1973 (2011) addressed to the President of the Security Council, UN Doc S/2021/229 (8 March 2021) para 63 <https://digitallibrary.un.org/record/3905159> accessed 26 August 2026.

¹⁶⁶*ibid*, annex 30, which is an identification panel recording the first sighting of the munition in the country, stating that its presence is a non-compliance with paragraph 9 of Security Council resolution 1970 (2011), and giving as its own primary sources a confidential source, the manufacturer's product page and a subscription reference service. On the reading of paragraph 63 see Hitoshi Nasu, ‘The Kargu-2 Autonomous Attack Drone: Legal and Ethical Dimensions’ (*Lieber Institute*, 10 June 2021) <https://lieber.westpoint.edu/kargu-2-autonomous-attack-drone-legal-ethical/> accessed 26 August 2026, observing that the report does not say for certain that human beings were killed by such systems operating without human supervision; and Zachary Kallenborn, ‘If a killer robot were used, would we know?’ (*Bulletin of the Atomic Scientists*, 4 June 2021) <https://thebulletin.org/2021/06/if-a-killer-robot-were-used-would-we-know/> accessed 26 August 2026. The manufacturer has stated that the system is not designed to launch fully autonomous attacks.

¹⁶⁷Final document of the Twenty-Seventh Annual Conference of the High Contracting Parties to Amended Protocol II, Geneva, 11 November 2025, UN Doc CCW/AP.II/CONF.27/5. The compilation searched the document for the terms recorded in the text and located none; the conference's substantive agenda concerned improvised explosive devices, technological cooperation under article 11, universalization and the article 13(4) national reporting synopsis. This is a statement about the contents of one document.

51. What a rule drafted now would reach

This paper takes no position on what the law should require. It has a narrow and unusual standing to say what the text now on the table *reaches*, because it has priced the mechanisms that text discusses, and it states that and stops.

The set of elements now before the Seventh Review Conference, in the report the Group of Governmental Experts adopted on 4 September 2026 at the close of its mandate, provides at paragraph 35 that the relevant measures of control and human judgement include, at sub-paragraph (c), “Measures to intervene in the operation of LAWS, to the extent possible or feasible, including to deactivate or neutralize LAWS in a timely manner”.¹⁶⁸ The rolling text of 18 December 2025 had provided at section III sub-paragraph 7(D): “Ensure that LAWS can be deactivated in a timely manner **and/or** incorporate self-destruct, self-deactivation or self-neutralization mechanisms”.¹⁶⁹

Two routes were offered in the earlier text and one survives; the 2026 drafts carried it as “That LAWS can be deactivated or neutralized in a timely manner”, the adopted text recasts it as a measure of intervention qualified “to the extent possible or feasible”, and the dropped route was not restored. The route that survives is the one requiring that something reach the system. The route that was dropped is the one that requires nothing to reach it.

On this paper’s own figures those two routes are not comparable in cost or in reliability, and the difference runs the opposite way to the drafting.

The surviving route depends on a command arriving. §25 prices the receive channel at about one part in seventy of the platform’s energy budget and is emphatic that this is cheap; but §25’s limits are equally emphatic that a receiver can be denied by jamming without being located, that the operator cannot know a command was received, and §27.1 states the consequence plainly—jamming a channel that carries a command to *stop* denies recall while leaving function untouched. **That consequence is a property of the direction the channel points, not of receive-only operation**, and §25.3 sets out the inversion that reverses it. A measure to “deactivate or neutralize LAWS in a timely manner” is, for an emplaced system under a competent electronic warfare effort, one satisfiable on a test range and defeatable in the field by an adversary who does nothing but transmit—though Appendix A.9 now bounds how much transmitting: on the processing gain a few-byte message permits, the jammer must sit within about a kilometre of the platform and stay there.

¹⁶⁸Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, Report of the 2024-2025-2026 sessions, UN Doc CCW/GGE.1/2026/3 (advance copy, 14 September 2026) para 35(c) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/CCW-GGE.1-2026-3-Advance_copy.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/CCW-GGE.1-2026-3-Advance_copy.pdf) accessed 15 October 2026. The Group adopted the report on 4 September 2026 (para 46). The draft of 7 July 2026, UN Doc CCW/GGE.1/2026/CRP.1, carried the provision at para 36(d) as “That LAWS can be deactivated or neutralized in a timely manner”, and the draft of 3 September 2026, UN Doc CCW/GGE.1/2026/CRP.1/Rev.1, at para 37(e) in the same words; v2.12.3 of this paper cited the July draft. This paper cites the text as adopted; what can still change is the decision of the Review Conference in November, which is a different question and is treated as such.

¹⁶⁹Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, rolling text, status 18 December 2025, section III para 7(D) [https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_\(2026\)/CCW_GGE_LAWS_Rolling_Text_-_status_18_December_2025.pdf](https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-_Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2026)/CCW_GGE_LAWS_Rolling_Text_-_status_18_December_2025.pdf) accessed 26 August 2026. Emphasis on the disjunctive is added. The chair’s summary of the first 2026 session records that some delegations supported deactivation and self-destruction measures as safeguards while others warned that they are not sufficient by themselves and should not substitute for human oversight: UN Doc CCW/GGE.1/2026/WP.2 (1 April 2026) para 40.

Two further things follow from that appendix and both bear on the phrase “in a timely manner”, which is doing more work in the text than its four words suggest. **Timeliness has a floor, and the floor is not the listening interval alone.** At the worked command-to-stop design point, the worst case without jamming is 60 seconds to reach a receive window **plus** 200 seconds of payload transmission at the rate its jam resistance requires, before acquisition, coding and implementation overhead: **at least 260 seconds.** A platform that listens hourly to conserve energy cannot be stopped inside the hour and three minutes, however good the channel. **Under a permission-to-continue design the guarantee is a different kind of thing:** not delivery of a stop command, but local safing after the authorised heartbeat interval and the message-validity allowance expire—which holds whether or not anything arrives. And **the schedule is an energy decision:** §21.2 prices the receive chain at one part in seventy of the budget precisely because it is duty-cycled, so a stricter timeliness requirement is purchased in watt-hours at a rate this paper can state. **And the schedule is not the only term.** §25.2 shows that the message itself takes two hundred seconds at the rate its jam resistance requires, so timeliness is bought twice over—once in listening energy and once in processing gain surrendered—and the second purchase is the one no amount of battery will buy back. A text that requires timely deactivation without specifying an interval leaves the interval to be set by whoever is paying for the battery. That is an engineering observation about a drafting choice, and the paper offers it as one.

A second drafting observation follows from §15.5 and is offered on the same footing. Wording that names a sensing modality is fragile, because a function can survive substitution of the component that performs it: a constraint written against camera-based recognition does not reach the same recognition performed on a point cloud, and one written against lidar-equipped systems does not reach it performed on an image. §41 states the general proposition and §49 the historical form of it. **The engineering observation is only that a functional description is more durable than a component description;** whether a treaty text should be drafted that way is a question for drafters, and the authority non-expansion invariant at §15.5.4 is a design constraint this paper proposes, not an obligation it claims to find in any instrument.

The dropped route depends on nothing arriving. §21.2 prices self-neutralization—an incorporated mechanism rendering the platform inoperable—at approximately nothing, because it is a property the architecture already has; self-destruction likewise; and §26 sets out why an independent second mechanism on a different physical principle, bounding the residue, is the better specification. §26 also records where that specification comes from: it is what Amended Protocol II has required of a remotely delivered mine since 1996, at ninety per cent reliability with a back-up leaving no more than one in a thousand functional after 120 days.

The comparison can therefore be stated in one sentence, and it is the plainest statement this paper has to offer the delegations to which its companions were sent. **On the text as it now stands, an autonomous weapon system would be asked for a capability that depends on a link an adversary can deny, and would not be asked for the mechanism that works when the link is gone—which is the mechanism the same body has required of a device with no sensing at all since 1996, and which this paper’s arithmetic prices at approximately nothing.**

Three further observations bear on how that should be read, and the paper offers them as observations rather than as recommendations.

The 2026 revisions also removed the sub-paragraph which would have limited such systems to “a defined perimeter” and restricted their use “to military objects”, and the adopted text does not

restore it.¹⁷⁰ §46 records that the target boundary—materiel rather than persons—is the one point on which a defence ministry’s own directive and the International Committee of the Red Cross converge. It is not in the current text.

The obligations of Amended Protocol II have become the operative floor for a group of states for whom they were previously redundant. Five European states deposited instruments of withdrawal from the anti-personnel mine convention between June 2025 and August 2025, with effect between December 2025 and February 2026.¹⁷¹ All five remain High Contracting Parties to Amended Protocol II. For them the Technical Annex regime—self-destruction, back-up self-deactivation, detectability—is now the binding constraint rather than a redundancy beneath a prohibition, and §21.2 prices every one of those obligations for the energy-spending case.

And the timetable is short, which bears on how this section should be read. **The text analysed here is the report the Group of Governmental Experts adopted on 4 September 2026 at the close of its mandate, issued as CCW/GGE.1/2026/3 in an advance copy of 14 September 2026.** Version 2.12.3 of this paper analysed the draft of 7 July 2026. The provision at the centre of this section passed from that draft to the draft of 3 September 2026 word for word and was then recast in the text adopted on 4 September: what the drafts stated as a capability, “That LAWS can be deactivated or neutralized in a timely manner”, the adopted text states as “measures to intervene in the operation of LAWS, to the extent possible or feasible, including to deactivate or neutralize LAWS in a timely manner”. The recast narrows the requirement further in the direction this section describes—a measure qualified by feasibility is satisfied by less than a capability—and the incorporated mechanism was not restored. The passages that contrast the December 2025 and 2026 formulations, at §25.3, §26, §29, §48 and the notes to Parts VII and IX, have been re-read against the adopted text, and the document number, the paragraph numbers, the quotations and the access date have been updated together. The advance copy is issued without formal editing; the check that remains is to read this section against the edited document when it appears. The group of governmental experts held the final session of its mandate from 31 August to 4 September 2026 and the preparatory committee for the review conference met from 7 to 9 September; the Seventh Review Conference sits from 16 to 20 November 2026, with the Amended Protocol II

¹⁷⁰Rolling text, status 18 December 2025 (n 169), section III para 7(E), providing for limitation of operation to a defined perimeter and restriction of use to military objects. No corresponding provision appears in the set of elements as adopted, CCW/GGE.1/2026/3 (n 168) para 35, whose sub-paragraphs run (a) to (c), or elsewhere in that text; none appeared in the drafts of 7 July or 3 September 2026 either. Paragraph 35(b) of the adopted text lists, as a measure, “Limiting the scale of the operation, types of targets, duration, and geographical scope, as appropriate”, which is a limitation of effects rather than a restriction of target class; the draft of 7 July 2026 had also listed “the number of engagements that LAWS can undertake autonomously”, which neither the draft of 3 September nor the adopted text carries. The chair’s summary of the first 2026 session records that many delegations wanted the restriction expressed as “military objectives” rather than “military objects”, while a number of delegations welcomed the limitation to military objects, stressing that there should be a prohibition on anti-personnel systems: CCW/GGE.1/2026/WP.2 (1 April 2026) para 41.

¹⁷¹Estonia, Latvia and Lithuania deposited instruments of withdrawal from the Convention on the Prohibition of the Use, Stockpiling, Production and Transfer of Anti-Personnel Mines and on their Destruction on 27 June 2025, with effect from 27 December 2025; Finland’s withdrawal took effect on 10 January 2026; Poland deposited its instrument on 20 August 2025, with effect from 20 February 2026, per depositary notification C.N.421.2025.TREATIES-XXVI.5 <https://treaties.un.org/doc/Publication/CN/2025/CN.421.2025-Eng.pdf> accessed 26 August 2026. The effective dates for the first four are taken from the reporting of the International Campaign to Ban Landmines and are not depositary notifications. Ukraine notified a decision to suspend the operation of the Convention in 2025; the Convention contains no suspension clause and article 20(3) provides that withdrawal does not take effect during an armed conflict in which the withdrawing party is engaged, and Ukraine’s status is not settled here.

group of experts on 10 and 11 November and its annual conference on 13 November.¹⁷² The report recommends that High Contracting Parties at the review conference consider and endorse the set of elements, and that the review conference decide by consensus on the possible modalities of the Group’s future work; it does not recommend the opening of negotiations. On 25 August 2026 the Secretary-General of the United Nations and the President of the International Committee of the Red Cross renewed a joint call for a legally binding instrument, stating that the review conference “offers the clearest path available”.¹⁷³

What this paper contributes to that is a number and not a view. The mechanisms in question have a price, the price is in the table at §21.2, and on the arithmetic there the cheapest of them is the one the current text no longer names.

Part X—The Boundary, and Conclusions

52. Limits of closure

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 §8 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.

¹⁷²Meeting of the High Contracting Parties to the Convention on Certain Conventional Weapons, final document, UN Doc CCW/MSP/2025/8, para 17(g), fixing the 2026 calendar including the Seventh Review Conference from 16 to 20 November 2026, the preparatory committee from 7 to 9 September 2026, the Amended Protocol II group of experts on 10 and 11 November 2026 and its annual conference on 13 November 2026; and CCW/MSP/2025/8 para 19(c) for the sessions of the Group of Governmental Experts. On the recommendations see CCW/GGE.1/2026/3 (n 168) para 45.

¹⁷³Renewed call of 25 August 2026 (n 163), on the review conference as the clearest path available.

The boundary of the argument is set by target speed. Ukraine’s Sky Sentinel marks where this paper’s architecture stops.¹⁷⁴ It is radar-cued, and it is radar-cued because its targets close at high speed: the speed of the target sets the sensing modality, the modality sets the emission posture and the power budget, and no closed, low-emission posture of the kind specified in Part V survives that requirement. Everything this paper proposes is available for ground denial because ground targets are slow. The capability reporting on the system traces to a single state-media article from which subsequent coverage derives, and a later report describes the same turrets as remotely controlled with an operator confirming the order to fire; the example is used here for its sensing and its target speed, which the reporting supports, and not for any claim about how the fire decision is taken. On the other side of that boundary, a dismounted man at walking pace can be detected at thirty metres by a passive geophone and engaged over tens of seconds.

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. Part IX places it among the wider record of fielded sentries and their modes; what matters here is only the trade it demonstrates. 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.¹⁷⁵ 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 require-

¹⁷⁴The capability reporting on Sky Sentinel traces to a single Ukrainian state-media article, United24 Media, 26 May 2025, from which David 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, and *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, both derive, as does the US Army lessons-learned paper at (n 121). No manufacturer is named in that chain, so the cost, performance and autonomy figures it carries have no identified originator. A later report, *Militarnyi*, 11 April 2026, names Carmine Sky as deploying the turrets and describes them as remotely controlled, with an operator confirming the order to fire. The example is used in the text for its sensing modality and its target speed, which the reporting supports, and for nothing else.

¹⁷⁵On 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 secondary accounts describe the programme as classified without citing a source, which is recorded here as a fact about those accounts; the figures here establish the system’s character rather than the scale of its use.

ment, 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 §28 does not resolve into a mitigation, and the paper does not pretend otherwise.

53. The state transition within the emplacement

A platform of the kind this paper describes spends almost the whole of its emplacement as a sensing device and almost none of it as a weapon: on the paper’s own budget, twenty-nine days of sensing against a firing event measured in hundreds of milliseconds. The question of when a machine of this class becomes an autonomous weapon is therefore not a question about the machine but about an interval within its operation, and a rule which classifies the machine as a whole will classify it wrongly for almost all of the time it exists.

Three limits hold this observation in place, and each is a limit rather than a qualification. The first: this is not an argument that such a system is therefore less regulated, or regulated only during the interval. The capacity to make the transition is present throughout the emplacement, and it is that capacity, not its exercise, which the paper argues a rule must reach.

The second: the proportion is a property of this worked case at its stated event rate, not a general property of the class. At a higher event rate the proportion changes, and the paper’s own sensitivity figures show the budget changing with it. Nothing in the observation survives being generalised to the class.

The third: the transition is not observable from outside. Nothing distinguishes a platform which is about to make it from one which is not, which is why the observation cannot be turned into a verification or inspection scheme, and the paper does not propose one.

One formal statement belongs here because a companion paper turns on it. Let t_e be the time of effect, τ_H the time of the last human decision that sets or enlarges authority, and t_m the time of an automatic modality change. The interval that matters is

$$I = t_e - \tau_H$$

and unless a person makes a new authority decision at the modality change,

$$I \neq t_e - t_m.$$

A camera-to-lidar transition may materially alter the evidence, the detection envelope and the identification envelope while leaving the human-decision interval untouched. The machine’s state changed; nobody decided anything. That is the same distinction §53 draws between a sensing state and a weapon state, applied to the sensing path rather than to the effector, and it is why an automatic transition cannot restart the clock: the clock measures human authority, and no human act occurred. The consequence for the specification is at §15.5.4, where the authority ceiling is fixed at τ_H and every modality’s usable envelope is required to sit inside it.

One concession closes the passage. The observation is a way of describing a design, not a discovery about it. A designer who set out to build such a system would arrive at the same proportion

without ever framing it this way, and nothing in the paper’s argument requires the framing to be adopted.

54. 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, 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; the quoted figure is therefore not an upper bound on emplaced duration but a measurement of a different problem. 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 a sourced sixteen hours of internal-battery endurance for a man-portable ground surveillance radar of the same class and era: the same task, a factor of about forty-five, and architecture rather than chemistry as the difference. The two figures come from different documents about different systems and are not a controlled comparison, which §17.2 states at the point of use. The calculation at §17.3 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 3 to 20 W spread assembled at §17.3 from separate manufacturers' figures for an uncooled imager and an edge inference module, the thirty-day requirement moves only from 13 Wh to 70 Wh, and the harvesting surface from 1.2 Wp to 6.3 Wp. Sampling every stated parameter simultaneously puts the ninetieth-percentile interval on the total at 0.80 to 3.59 Wh per emplaced-day, and the discrimination multiple never falls below about six and a half anywhere in it. The design is insensitive to a factor of nearly 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 89 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 1.4 per cent, and the weapon itself at 0.4 per cent. A mine is the sensorless limit of the same design—it performs no sensing 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, mechanism by mechanism: self-neutralization, which is what the architecture supplies, and self-destruction are cheap; detectability costs approximately nothing; self-deactivation, which the instrument defines as the irreversible exhaustion of an essential component, is in tension with an architecture built to prevent exactly that exhaustion and must be supplied by a separate component; 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 continuous lidar ranging out of both the emission and the power budget and is part of why §17.3 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. The survey-integrity interlock of §15.4 answers it: a landmark set surveyed once, a robust pose fit against it on each wake, an always-live inertial channel, and a release condition *V* that expires unless both agree the platform still occupies its surveyed frame. The condition is a conjunction, so an obscured scene—which cannot be evaluated—expires it as surely as a displacement does, and ambiguity therefore safes. Both channels are closed on the test at §6.1, an inertial unit as much as the camera; what an inertial unit cannot do is navigate over thirty days, and what it is used for here is to detect an event. 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 §27 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 seventy 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, $N_3 = Tp_1p_2 + F_1q_2$, 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 *command-to-stop* 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. On the permission-to-continue channel of §25.3 the same jamming safes the platform within one authorisation period, so the perverse incentive is a property of the message’s sense rather than of receive-only operation. 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.

The transition from sensing to effect is an interval within the emplacement, not a property of the machine. On the paper’s own budget such a platform spends twenty-nine days sensing against a firing event of hundreds of milliseconds, so the question of when it becomes an autonomous weapon is a question about an interval within its operation and not about the machine, and a rule that classifies the machine as a whole classifies it wrongly for almost all of the time it exists. The point sets three limits and no more: the capacity to make the transition is present throughout and is what a rule must reach; the proportion holds only for this worked case at its stated event rate; and the transition is not observable from outside, so it supports no inspection scheme.

The public record establishes that the capability class exists, is funded and is physically calculable, and it establishes nothing about performance. Uncrewed ground platforms, remote weapon stations, the fire control that aims them, target recognition and photon-counting optical sensing are separately named, separately funded and separately contracted in published budget exhibits and the federal award record, with entries current to 2026; the award record states performer, amount and period of performance, and does not state what was delivered. The physics that determines what an optical sensing and fire-control architecture can do is open literature, most of it decades old, and a bound on what no hardware can beat is available to a rule-maker without knowledge of any particular system. What is not public is which arrangements a given party has built and how well they work; the record marks its own boundaries, and a constraint drafted against named programmes or named technologies depends on a naming that the record shows to be unstable, where a constraint drafted against function and against use does not.

The safeguard must be architecturally independent of the sensing it restrains, and the commodity market prices the same inversion. In fielded machine-vision weapon stations the human approval is a layer on an autonomous sensing-and-aiming core, so the safeguard shares a substrate with the targeting and is a single point of failure; this specification places release on a channel that shares nothing with perception, behind physical safing and an expiry, which is what makes it a property of the machine rather than a setting. And the components outside the discriminating tier—platform, sensor, effector—are listed to the public at consumer prices, so the device that does not discriminate is the simpler and cheaper one, which is why discrimination has to be required rather than assumed.

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. §20.3 states the experiment.

55. Sensitivity of the conclusions

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 the 111 mWh assumed at §17.3. The exchange rate at §21 rests on that figure, and the whole reframing of energy as the price of discrimination moves with it; §17.4 shows the proposition survives the stated uncertainty, but a measured figure an order of magnitude higher would not be within 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.

56. Summative remarks

This paper started from the published record: endurance for these platforms is quoted in drive hours, and no published figure describes what a system that does not move can do.

There is no such figure because the field measures drive hours, 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.

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 §28, 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.

Part VIII adds one thing, and it should be said in ordinary words because the audience the companion papers address will need it in ordinary words. Everything the specified object requires already exists; most of it is cheap; its components are offered for civilian sale in the United States, subject to the jurisdictional, export and end-use restrictions that attach to them; and in the systems now sold and fielded the safeguard is a layer that can be removed. The two published companion papers state that the object they analyse is not hypothetical in its components. This paper is the record that makes that statement checkable, and §41 states what the record does and does not support.

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.

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—Validation Requirements

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 principal experiment

The event-rate measurement. The budget at Part VI §17.3 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.

Label every logged event with the sensing modality and configuration that produced it, and carry ground truth sufficient to estimate each cascade term on its own rather than in combination. *Measure:* not “escalations per day”, which does not transfer between sites, but the three terms of the cascade at §17.4 **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 §17.3 is realistic. Two hundred tier-one trips a day give ten escalations and 1.22 Wh per emplaced-day. Eight hundred trips give 4.04 Wh, which the interlock at §18 still supports. Two thousand give 9.67 Wh, and the harvesting surface has grown beyond the size the concealment argument at §18.1 assumes. The paper’s arithmetic assumes ten and the assumption is not defended by evidence, only by plausibility.

Report P_D , $P_{ID|D}$, F_1 and q_2 with sample sizes and confidence intervals, and do not report accuracy alone: the event set is heavily imbalanced, and an accuracy figure on it is uninformative about exactly the term that dominates the budget. 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: measure a surface, not a floor. The design allows 500 μW for the always-live tier at a 30 m detection range. §17.3 now anchors the electronic order of magnitude of the acoustic part of that allowance against a published 305.5 μW wake-up chip and shows that doubling the whole allowance moves the total budget by under one per cent. **What remains unsourced is the pairing,** not the wattage: no source establishes that an integrated seismic-and-acoustic tier detects a walking person at 30 m at any stated probability, for any stated false-trip rate, at that power.

Measuring the power alone would close half of it. **The measurement worth making instead produces the whole surface §17.5 needs**, because the same instrument run yields it at no extra cost: for an integrated seismic and acoustic front end with wake-on-event—microcontroller, real-time clock and conditioning included, **and those three charged to P_1 and therefore excluded from P_{hk} , or measured within P_{hk} and excluded from P_1 , so that a combined idle-rail reading is charged once and not twice (§17.3)**—sweep the detection threshold and, at each setting, record

- the continuous draw of the tier;
- the probability of detection against range, **with a stated confidence interval and a stated sample size**, for a walking person, on ground representative of the intended siting and across the band the front end actually accepts—reported as P_D and not as an accuracy figure, since the event set is heavily imbalanced and accuracy on it is uninformative;
- the false-trip rate per day, over a period long enough to include weather and wildlife;
- and, derived from the last of these through the cascade at §17.3, the daily energy and the thirty-day requirement.

The output is a Pareto frontier over threshold: validated detection range against false trips against watt-hours per emplaced-day. That is the figure §17.5 draws from stipulated relations, measured. **It converts the energy-constrained detection range from a framework into a result**, closes the 500 μ W assumption and the F_1 assumption in the same experiment, and requires no weapon, no release mechanism, no autonomy and no controlled equipment. *Decides*: whether the always-on floor is where the unattended-sensor literature implies it is once a detection requirement is imposed on it, and what a sector actually costs to watch.

The parameters that experiment does not close should be listed with it, because §17.3's vector is not all of the budget: the tier-two component class of 50 mW and its ten-second dwell, the 2 mW housekeeping load, the tier-two leakage q_2 , the genuine-event rate T , the receiver's 30 mW and its message frequency, and the engagement count are stipulations or component-class figures rather than measurements, and are labelled as such where they are used. A.8 states the whole set. The tier-one surface is the one whose absence most weakens the paper, which is why it is first.

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 prices the corrected receive chain at 17.0 mWh per day, on a 30 mW scheduled receiver and the two-hundred-bit authenticated message of §25.2. 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 §15 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 §30 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 0.020 m^2 , 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 §18.1 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 §18.2 holds, or whether an uncharged buffer produces a failure rather than a lengthened response interval.

A.4 The scavenging feasibility study

Stated at §20.3 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 §20.2 in place from the outset rather than added after an incident.

A.5 The survey-integrity test

Part V §15.4 makes survey integrity a release interlock and states it as the condition $V(t)$. That condition is the paper's strongest safety proposal and it is entirely untested. This experiment tests it, and needs no weapon, no release mechanism and no armed platform.

Build: an instrumented but unarmed platform carrying the tier-three imaging chain, a low-power inertial unit, and the pose-fit routine of §15.4—survey a landmark set at emplacement, re-observe it on each wake, solve for the pose that best explains it with robust outlier rejection, and evaluate the four clauses of V .

Do: subject it to a schedule of **controlled displacements**—translations and rotations of known magnitude, from a few centimetres upward, applied by hand, by vehicle contact and by a shock table—and, separately, to **controlled scene alterations**: removing, adding, occluding and moving landmarks while the platform itself is undisturbed. Run both across scene types from feature-rich to near-featureless, and across light, season, snow cover and partial obscuration.

Measure: translation and rotation separately, because the corrected condition at §15.4 treats them separately and they have different consequences at the boundary. For each trial record the inlier count m and residual ρ ; the translation and rotation estimation errors against the applied ground truth; the probability that V survives a displacement whose boundary effect exceeds b^* ; the nuisance safing rate with no displacement applied; and the freshness-failure rate. Where the hardware permits, report inertial-only, vision-only, lidar-only and fused results separately, since §15.4 claims the pairing is strictly better than either channel alone and that claim is falsified if one channel catches everything the other does.

Error	What it means	Why it matters
V survives a displacement	the platform believes it still occupies a frame it has left	the safety-relevant failure ; a specification should bound it
V expires without a displacement	the platform safes itself for a benign reason	the availability failure; sets whether the interlock is usable

Also record the smallest displacement reliably caught by each channel separately—inertial and visual—since the claim at §15.4 is that the pairing is strictly better than either alone, and that claim is falsified if one channel catches everything the other does.

Decides: whether the geometric authority constraint, which is the foundation of the whole authority envelope at §15.6, can be relied on at all. If V survives displacements it should not, the constraint is decorative and the specification is unsound. If V expires constantly under ordinary weather, the interlock is correct and unusable, and the design must fall back to the inertial channel with a human clearing every alarm—which is a different and much more expensive concept of operations, and one the energy argument does not price.

A.6 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.7 What the paper got wrong in draft, recorded because it bears on the argument

The corrections below were made during preparation and revision. Each is recorded because it bears on the argument, the audit trail or both.

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.

The paper opened on a reported case of an armed ground platform holding a fighting position for over a month, treated the case as evidence that quoted endurance figures measure the wrong thing, and described the thirty-day design case as having been arrived at from that report and as converging independently on the same figure the treaty fixes. On the published reporting, the platform was under continuous remote control from an operator about ten kilometres away, with uncrewed aircraft overhead supplying targeting in real time; a second account of what appears to be the same deployment reports that the operators took the platform to the position each morning and returned it each evening. It was neither unattended nor offline nor continuously emplaced, and a platform recovered daily is entirely consistent with the endurance figures the paper called into question. The account containing the daily-recovery fact is the article the paper cited at its own first footnote in v1.0.0; the correction was inside the source already relied on. The opening no longer rests on the case. The endurance argument re-anchors on a comparison the paper already carried, between a fielded thirty-day unattended sensor system and an eight-hour portable surveillance radar, both published and both unaffected; the thirty-day design case is now stated as what it is, a stipulation taken from the treaty; and the claim of independent convergence is withdrawn. The error reaches beyond this paper: a companion paper, already published, describes a platform "of substantially that description" and refers to a horizon "arrived at independently from an operational case", and neither statement is supported; both are corrected in an annex circulated to the delegations to which those papers were sent. Version 2 of each companion paper (15 October 2026) adds a postscript on the report the Group of Governmental Experts adopted on 4 September 2026 and leaves the body unchanged. The general lesson is narrower than the error. A source cited for one proposition was not read to the end for the others it carried, and the paper's most quotable sentence was the one least examined.

The paper priced one of the instrument's three obligations at approximately nothing, on the ground that it was a timer, and concluded that it and a second obligation were both close to free for the architecture specified. The instrument defines that obligation not as a timer but as the irreversible exhaustion of a component essential to the device's operation; the paper had matched the name to the wrong mechanism. What the architecture actually supplies is a different one of the three—an incorporated mechanism that renders the device inoperable—which the paper had not named at all. The error was not terminological. On the correct definition the obligation is the one the paper's own central contribution works hardest against, since the whole energy argument is about preventing a component from irreversibly exhausting, and the paper had priced at approximately nothing the requirement its architecture is worst placed to meet. Each of the three mechanisms is now named and priced separately; the price of the misnamed one changes sign, from free to in tension with the architecture; and the honest engineering statement is added, which is that the function has to be supplied by a separate exhausting component rather than by

the main energy path. The lesson is that the paper took a legal term from the secondary literature rather than from the instrument, and the secondary literature uses it loosely.

The paper claimed priority for a doctrinal idea that a fielded programme had already implemented in part. §30 stated that a closed holding system at posture 1 is “the first obstacle whose intent can be changed after emplacement without anyone going forward”. An emplaced munition whose arming state can be commanded from the rear has been fielded since at least the modular pack mine system, and the M7 Spider and the United States Army’s terrain-shaping obstacle effort are described in the Army’s own publications as carrying networked two-way command with on-off-on and remote-destruct capability, for the stated purpose of keeping emplaced effects under friendly control. The claim of priority is withdrawn. What replaces it is narrower and survives the record: those systems make *arming* mutable, and the object specified here makes *target and obstacle effect* mutable, because it carries a discriminating sensor whose sector and release criteria are parameters rather than a detonation state. The error was one of literature search rather than of reasoning—the prior art is in the open publications of the same army whose doctrine the section cites—and it was found by searching the military record rather than the academic one.

The paper stated that discrimination “was not available at any price” for emplaced devices. That is true only of discrimination in the sense the law of targeting requires, between a person who may be attacked and one who may not. Discrimination between classes of materiel by signature was fielded in emplaced munitions from the 1970s, and §21.1 now says so and distinguishes the two senses. The correction narrows the claim and makes the price at §21.2 a price for something the record shows was genuinely missing, rather than for something the record shows had been done.

The recall channel was priced before its link was computed, and the price was wrong by an order of magnitude. Versions to v2.8.0 costed the receive chain at 0.9 mWh per day on a two-stage architecture whose always-listening stage has a sensitivity of about -90 dBm. Appendix A.9, added at v2.8.0, then computed the link over a two-ray plane-earth model and found that such a receiver is not reached by a twenty-watt station beyond about two kilometres, against a design case of five to fifteen. The channel that was priced does not close at the range the specification states. §25.1 now prices the channel that does—a sensitive receiver held up on a schedule—at 17.0 mWh per day, 1.4 per cent of the budget, and the wake-up stage is retained for the close-range path it can actually serve. The lesson is narrower than A.7’s usual one and worth stating for that reason: **the paper computed the power of a component before it computed whether the component could do its job, and the two calculations were done a version apart.**

The recall message was assumed to be sixty-four bits and one second, and both were wrong. §25.2 now specifies two hundred bits—address, command, anti-replay counter and authentication tag—and shows that delivering them at the rate which yields sixty decibels of processing gain takes two hundred seconds. A one-second window detects the carrier; it does not deliver the message. Worst-case recall latency at the design point is therefore about four minutes rather than one, and the correction produces the coupling law at §25.2 that the paper did not previously have.

The negative finding on documented autonomy was overtaken, and the evidentiary claim under it was wrong. §50 recorded that no independently documented autonomous engagement had been located, and reasoned that none could be documented because an autonomous and a commanded engagement are not distinguishable from the wreckage. On 24 August 2026 a case was published in which they were distinguished from the wreckage, because the compute module was recovered and its software was unencrypted. The negative is withdrawn and narrowed to autonomous selection of *persons*; the evidentiary claim is replaced by the corrected one, that

autonomy is not visible in the effect but is visible in the artefact. §50 sets out both, and the consequence for a verification regime is now stated there rather than merely conceded.

Three legal statements were wider than their instruments. Protocol IV’s precaution obligation was said to attach to “any architecture that emits optically”, which is not what Article 2 says: it attaches to the employment of laser systems, and a passive optical stage is not one (§15.2). §46 was headed “already lawful”, which a national policy directive cannot establish and which its own text does not claim; the heading and the opening now say what the directive is, which is a policy exemption from a senior review requirement (§46). And Amended Protocol II’s thirty-day and 120-day figures were attributed to remotely delivered mines generally; they attach to remotely delivered *anti-personnel* mines, and to non-remotely-delivered anti-personnel mines outside a marked perimeter, while other remotely delivered mines carry only a “to the extent feasible” obligation with no figures at all (§21.2). The last of these strengthens §51 rather than weakening it.

The displacement argument mislabelled a closed sensor as an open one, and asserted an invariant that is not invariant. §15.4 preferred visual self-survey over an inertial unit on the ground that only the former passes the withdrawal test. An inertial measurement unit depends on no third party and passes it as cleanly as anything in the specification; the objection to it is drift, which disqualifies it from navigation and not from detection. The same subsection stated that inter-feature geometry is invariant under displacement of the observer: included angles are invariant under rotation but not under translation, which is the displacement that matters. Both are corrected, and the corrected design pairs an always-live inertial trigger with a visual landmark fit that is permitted to return *ambiguous*. Misapplying this paper’s own central distinction is the more embarrassing of the two errors and is recorded first for that reason.

A seismic propagation statement conflated amplitude with energy, reproducing an inconsistency in its own source. The text gave body waves as losing 6 dB per doubling while the footnote gave them as falling with the inverse square of range; the first is an amplitude statement and the second an energy one. The cited study contains both forms, and its own conservation-of-energy derivation supports amplitude falling as the inverse of range for body waves and as the inverse square root for Rayleigh waves. §15.1’s note now states the quantity, and the conclusion—that the surface wave reaches further—is unchanged.

A flagged citation was resolved by withdrawing the figure rather than by finding it. Footnote 48 in v2.0.0 carried a bracketed note recording that the eight-hour endurance attributed to a ground surveillance radar could not be traced. It still cannot; no primary source for it was located. The comparison at §17.2 now rests on a sourced sixteen-hour internal-battery endurance for a man-portable ground surveillance radar of the same class, taken from a national forces manual, and the ratio is restated from about ninety to about forty-five. The paper also now says at that point that the two figures come from different documents about different systems and are not a controlled comparison. The lesson is the one A.7 keeps recording: the sentence that travelled furthest was the one whose source had been checked least.

The budget’s internal arithmetic did not close, and correcting it changed the headline figure. §17.3 costed tier two at thirty tier-one trips a day while §17.4’s cascade, and the tier-three line derived from it, assumed two hundred; §17.2 referred to the second figure while §17.3 used the first. The two cannot both be right, and the tier-two line was the error: tier two is woken by every tier-one trip, so at two hundred trips it costs 28.0 mWh a day rather than the 4.2 mWh printed. The presentation has been rebuilt so that the error is not available to be made again—every line

in §17.3 is now derived from a single stated parameter vector rather than assumed independently, and Appendix A.8 gives the procedure that reproduces every figure from it.

Two further corrections were made in the same pass. The budget assumed the inference tier drew its power only while live and said nothing about what holding it ready costs; a settle-and-resume term is now carried explicitly, and §17.3 states as a specification requirement that the tier be hard-gated, with the penalty for a retention state tabulated. And every figure was a point estimate; §17.4 now gives the interval within which the result sits when all ten varying parameters are sampled together.

The effect on the paper's conclusions should be stated plainly, because it is not uniform. The headline moves from 0.91 to **1.22 Wh per emplaced-day** and the thirty-day requirement from 27 Wh to 36.6 Wh—both still trivial against any platform's mass budget. The discrimination share moves from 91 to 91 per cent. The discrimination multiple, which the companion papers quote, was eleven and is **about eleven**, so that figure survives the correction rather than being undone by it. And the proposition the paper actually rests on—that identification and release judgement dominate the budget by about an order of magnitude—is now shown to hold across the whole of the stated uncertainty rather than at a single point, which is a stronger claim than the one it replaces.

The specification asserted a link it had not computed, and computing it changed the specification. §25 priced the recall receiver's power to three significant figures and left its link budget as an assertion; three readers identified that as the gravest omission in the paper. Appendix A.9 now computes it, over a two-ray plane-earth model rather than a free-space one, because a concealed antenna at ground level suffers 40 dB per decade of range and not 20—a difference of 50 dB at ten kilometres, which is a factor of one hundred thousand in required power. The result does not support what the paper implied. On the always-on wake-up stage's cited sensitivity, the control station at five to fifteen kilometres in the design case at §15 is unreachable at any transmit power a covert station would use. What closes the design case is the *gated* receiver in a scheduled listening window, which §15.9 already specified but which the paper had presented as a means of saving energy rather than as the mechanism on which recall depends. The correction is therefore not to the arithmetic but to what the arithmetic was for, and it produces a result the paper did not have: the latency of recall is bounded by the listening schedule and not by the channel, which is what §51 now says to the text before the treaty body. The jam budget in the same appendix runs the other way and strengthens §27.1, bounding the standoff within which a jammer must sit. The general lesson is the one this Appendix keeps recording: the quantity left uncomputed is the quantity that was wrong.

A body of fielded evidence was obtained late and it corroborated more than it corrected. Thirteen documents were supplied to this compilation after v2.5.0, several of them scanned government reports requiring optical character recognition. Four bear materially on Part VI and are now cited there. A manufacturer's product literature for a fielded seismic unattended ground sensor gives a one-month mission life on three AA primary cells and an imaging node woken only by a sensor detection, which is this paper's design case and its architecture as a catalogue item; a further paper from the same manufacturer gives the core system at approximately 10 mW against the 2.5 mW this paper assumes for its always-on floor. A 1998 thesis records that a change of environment alone moved a fielded sensor activation rate by a factor of three to four, and that the campaign's own remedy for false activations was to require detection in more than one modality before reporting—which is the cascade at §17.3, adopted under operational pressure fifty years earlier. The effect on the paper is that several assumptions became reported figures and one, the tier probability of 0.95, is now visibly at the optimistic end of what the current literature reports. Nothing in the arithmetic changed. What changed is how much of it rests on assertion.

A source was fetched under a wrong identifier and the error was the paper's. The list of sources to obtain gave a digital object identifier for a 1955 paper on the partition of energy between elastic waves. That identifier resolves to a different paper in the same volume, on hydromagnetic turbulence, and the document retrieved under it was accordingly of no use. The paper's own citation was wrong, not the retrieval. The claim it was to have supported is not affected, being already carried by a source read in full, and the intended reference is recorded here by title, volume and pages rather than by identifier so that it can be found by someone who wants it: Miller and Pursey, 'On the partition of energy between elastic waves in a semi-infinite solid', *Proceedings of the Royal Society of London A* volume 233, pages 55 to 69.

A range figure travelled through three versions attached to no document. §15.1 stated that seismic footstep detection extends to about 94 m under favourable conditions and that stealthy footfall had been evaluated at about 51.5 m, citing them to "related range-limitation work" alongside a named conference paper. The named paper has now been read in full and contains neither figure, the related work was never identified, and no document carrying either number could be found. Both are withdrawn. What replaces them is better than what they claimed: the same source gives seismic transmission loss as a geometric term plus an excess attenuation proportional to range *and frequency*, with a stated attenuation constant, which yields the design result that detection range depends on the passband as well as the ground—a stronger statement than the two withdrawn figures, and one that can be checked. The lesson is the one recorded twice already in this Appendix, and it has now cost the paper three separate claims: a figure quoted without its source read to the end is a figure that cannot be defended, however plausible it looks and however long it survives.

The paper reported several documents as inaccessible which were accessible by other means. Part VIII recorded that certain Army narrative justification volumes and certain small-business technical abstracts could not be consulted. On a second attempt using different access methods both were retrievable, and the Army volume yielded a fact now recorded at §40: programme element 0604531A appears in the FY2027 Budget Activity 04 volume's list of programme terminations. What could not be consulted at the time was the full text held by one optics conference publisher, whose platform interposes an automated-access control that this compilation did not attempt to circumvent. Two of those papers have since been supplied to the author from his own copies and are now examined in full; two more remain unobtained and nothing rests on them. **One further correction belongs with them: the paper previously grouped four Geiger-mode sources as inaccessible SPIE records. One of the four is an IEEE Aerospace Conference paper and not an SPIE one, and it has in any event been displaced from its load-bearing position by a 2012 SPIE camera paper now examined in full.** The general lesson is the one this Appendix keeps recording in different forms: a failure to retrieve is a fact about the attempt before it is a fact about the record, and the two should not be reported in the same words.

Eight source documents were obtained and examined, and four of the paper's statements changed as a result. The documents were supplied to the author, from copies he already held, after the paper had recorded that one optics publisher's automated-access control had prevented retrieval; the control was not circumvented, then or now. Three are the exact papers the paper had cited on bibliographic record alone—Marino and others (2003), Cova and others (1996) and Bijelic and others (2020)—and each is now examined in full and its evidence class upgraded. Four are different papers that support the propositions independently: Entwistle, Itzler and others

(2012), Aull and others (2018), Aull (2016) and Albota and others (2002). One is an industry article by the manufacturer of a current commercial system.

The four changed statements, each identified with the document that changed it. *First*, the foliage claim. **This entry is stated twice over, because v2.12.0 overcorrected and v2.12.2 corrects the overcorrection.** The paper had recorded partial foliage penetration as a bounded result in the literature, resting on two SPIE papers it had not obtained. Aull and others (2018) states that the capability was obtained by an airborne system “fusing multiple 3D images taken from different look angles”, and Marino and others (2003) records that a single viewing direction leaves the foliage shadow as holes in the scene and that the sensor is deliberately desensitised to see past near cover. v2.12.0 restated the claim as a multi-look capability an emplaced platform cannot execute, and the two unobtained papers now carry nothing. **That went too far.** Marino also records that objects under trees *were* measured from a single viewing direction; returns beneath a canopy arrive through gaps, and angular diversity raises the probability of an unoccluded path rather than being the sole means of obtaining any return at all. **v2.12.2 states the narrower proposition:** a fixed viewpoint may recover some under-canopy geometry and retains occluded volumes; the robust wide-area result reported for the airborne class depends on angular diversity and multi-look fusion and cannot be transferred to a fixed horizontal emplacement without measurement. The desensitisation is likewise restated as a trade at an operating point, which a design may buy round by interleaving settings at a cost in time and energy, rather than as an architectural impossibility. *Second*, the receive-chain load. Aull and others (2018) gives 0.5 W as typical dissipation for a framed readout chip alone, and a manufacturer’s page gives 18 W for an integrated camera; Marino and others (2003) gives an average optical output of about 0.33 W for the transmitter. **The paper had implicitly reasoned from the transmitter, and §15.5.1.2 corrects that:** single-photon sensitivity says nothing about the receive chain, whose dissipation does not scale down with pulse energy, and the cited readout operating point is a thousand times the tier-one allowance and therefore cannot be the continuously powered tier. **v2.12.0 went further than the evidence and said the energy sits in the receiver rather than the laser. It does not follow:** the transmitter figure is optical output and the receiver figures are electrical dissipation, no wall-plug efficiency is reported, and the two cannot be ordered. The claim is narrowed in v2.12.2 wherever it appeared. *Third*, the sweep cost. Marino and others (2003) Table 1 gives an instantaneous field of view of 10.3×10.3 mrad, from which §15.5.1.3 computes the number of pointings and the energy of one sector sweep; the paper had classified scheduled search as a different posture without pricing it. **v2.12.0 measured the sector as a planar product of two angles, which overstates a $60^\circ \times 10^\circ$ solid angle by 4.6 per cent, and described every input as a floor when the proxy powers are not lower bounds. v2.12.2 uses the exact rectangular solid angle—1,644 pointings rather than 1,723—and describes the result as a deliberately favourable illustrative estimate. The conclusion is unchanged: one sweep an hour costs more than the entire passive architecture.** *Fourth*, the eye-safety coupling. Aull and others (2018) Table I gives a fiftyfold dark count rate difference between silicon and InGaAsP/InP on cooled test devices, and Entwistle and others (2012) reports a higher dark count rate as the primary tradeoff at $1.5 \mu\text{m}$; §15.5.1.4 is new and connects a Protocol IV precaution to the false-alarm term of the energy budget.

Two smaller corrections belong with them. The paper had described four Geiger-mode sources as inaccessible SPIE records; one of the four is an IEEE Aerospace Conference paper, and it has in any event been displaced from its load-bearing position by a 2012 SPIE camera paper now examined in full. And the Part VIII compilation described Cova and others as cited and not read in full; it has now been read in full and the description is corrected.

A cross-reference defect in v2.11.0 is corrected, and it is recorded because of how it arose. The v2.11.0 build renumbered the subsections of §15 after inserting a new one, and the renumbering pass ran over prose written earlier in the same build, bumping six references to the new modality subsection from §15.5 to §15.6. Six other references to §15.6 were correct and are unchanged. No proposition was affected; a reader following the wrong references would have arrived at the authority envelope instead of the modality analysis. **The general lesson is a build lesson rather than an argument lesson: a mechanical pass that rewrites tokens must run before, not after, the passes that write new prose containing those tokens.**

Version 2.12.3 corrects a display-only regression introduced when the v2.12.2 master was assembled. The first row of §17.3’s tier-one robustness table had reverted from the correctly rounded 1.2167 Wh/day to 1.2166 Wh/day. The exact result is 1.216670 Wh/day ($1.221338 - 0.012 + 0.007332$), which rounds to 1.2167. The model, parameter vector, equations and all other tabulated outputs are unchanged; this was a transcription and truncation error, not a modelling change.

Version 2.12.4 updates the treaty analysis to the text the Group of Governmental Experts adopted on 4 September 2026. The Group of Governmental Experts adopted the report of its 2024-2025-2026 sessions on 4 September 2026 and it was issued as CCW/GGE.1/2026/3 in an advance copy of 14 September. The citations to the draft of 7 July 2026 and to the rolling text of 5 June 2026, at §4, §25.3, §26, §29, §48, §51 and the notes to Parts VII and IX, are replaced by citations to the adopted text. The finding at §51 is unchanged: the provision survived to the adopted text, at paragraph 35(c) in place of 36(d), recast as a measure of intervention qualified “to the extent possible or feasible”, and the dropped route was not restored. §15.5.5 gains a paragraph recording that the qualifier the drafts carried for the re-review of modified systems—“significant”, then “substantial”, then “material” impact—was dropped with the provision itself from the adopted text, and the companion note’s statement about the text before the Group is confirmed against it. The pre-publication caveat at §51 is replaced. No model input, figure, table or numerical result changes, and the reproducibility archive of v2.12.2 remains controlling.

Version 2.12.5 corrects statements that gave targets and plans as accomplished, and withdraws figures that the sources cited for them do not contain. The 25,000 unmanned ground vehicles for the first half of 2026 were a target announced by Ukraine’s Defence Minister, not vehicles contracted, and the 4.5 million first-person-view drones for 2025 were a planned procurement; the extended summary, §16.1 and §20.1 now say so, and the note to §16.1 records the target. At §20.1 the Ukrainian production figures for 2022 and 2025 and the monthly figure are withdrawn, because the source cited for them gives only the capacity for 2026, and so are the Russian comparison with the preceding year and monthly estimate, which the report cited for them does not contain; the extended summary now gives the combined figure as the two sides’ stated capacity and plans. The share of logistics performed by ground robots is attributed to the one brigade and two cities its source names, and dated; the capture of a position by unmanned platforms alone is described as the first reported, as the President’s statement has it, rather than as the single case, and is dated and placed; and the logistics fleet is described as robotic rather than autonomous, since its source describes vehicles controlled by operators. Two statements of the discrimination multiple are aligned with §17.4. The earlier-title note gives the subtitle of v1.0.0 and records that the legal companions refer to this paper without a title. The legal companions are now identified by their all-versions DOIs, with their version 2 recorded, and A.10 states where the reproducibility archive is published and gives its SHA-256. Seven cross-references that named the wrong note or section are corrected, and five statements about this or the preceding version of the paper that

had ceased to be accurate are corrected or removed. No model input, equation, figure, table or numerical result changes, and the reproducibility archive of v2.12.2 remains controlling.

A.8 Reproducing the budget

Part VI's figures are not transcribed; they are computed from the parameter vector at §17.3 by the procedure set out here, and any reader who disagrees with an assumption can substitute their own and recompute.

The vector. $P_1 = 500 \text{ } \mu\text{W}$; $T = 2/\text{day}$; $F_1 = 200/\text{day}$; $p_1 = p_2 = 0.95$; $q_2 = 0.041$; $P_2 = 50 \text{ mW}$; $t_2 = 10 \text{ s}$; $P_3 = 10 \text{ W}$; $t_3 = 30 \text{ s}$; $t_b = 10 \text{ s}$; housekeeping 2 mW ; receiver 30 mW held up for 1 s in every 60 and for 200 s on each of three messages a day, giving $12.0 + 5.0 = 17.0 \text{ mWh/day}$; effect $20 \text{ W} \times 5 \text{ s} \times 5 \text{ engagements}$ over a thirty-day horizon.

The derivation. $N_2 = T \cdot p_1 + F_1$. $N_3 = T \cdot p_1 \cdot p_2 + F_1 \cdot q_2$. Tier one is $P_1 \times 24 \text{ h}$. Tier two is $P_2 \times t_2 \times N_2$. Tier three is $P_3 \times (t_3 + t_b) \times N_3$, plus any retention power $\times 24 \text{ h}$. Housekeeping and the receiver are as stated. Effect is the engagement energy divided by the horizon. The total is their sum; the indiscriminating variant is the total less the tier-three line; the multiple is their ratio. The readiness interval is $t_2 + t_b + t_3$.

The sensitivity. Elasticities at §17.4 are computed by perturbing each parameter by one per cent and recording the proportional change in the total.

The interval. The distribution at §17.4 is obtained by sampling $P_3 \in [3, 20] \text{ W}$, $t_3 \in [20, 45] \text{ s}$, $t_b \in [5, 20] \text{ s}$, $F_1 \in [100, 400]/\text{day}$, $q_2 \in [0.020, 0.082]$, $T \in [1, 5]/\text{day}$, $P_1 \in [300, 800] \text{ } \mu\text{W}$, $P_2 \in [30, 80] \text{ mW}$, $t_2 \in [5, 15] \text{ s}$ and housekeeping $\in [1, 4] \text{ mW}$, each on a triangular distribution with its mode at the design point, over 200,000 draws. **The sampling support is a modelling choice and not measured uncertainty, and the result is a sensitivity ensemble rather than a confidence interval.** The triangular distribution encodes that the design point is the most likely value; the stated limits—including the $300\text{--}800 \text{ } \mu\text{W}$ on P_1 —are **illustrative support limits chosen for sensitivity analysis**, not empirically established hard bounds, because no measured distribution for any of these parameters exists and nothing rules out a value outside them. A reader who prefers uniform sampling will obtain a wider interval and the same conclusions.

The figures. Figures 1 and 2 are plots of the expressions above, evaluated over the ranges stated in them, and carry no data that is not in this appendix. **Figure 3 is the illustrative detection-envelope example of §17.5 and is a plot of stipulated sensing relations rather than of measurements.** Figure 4 is the recall-coupling figure of §25.2. Figures 5 to 7 stand in the same relation to the expressions at A.9.

The computation is elementary enough to be done in a spreadsheet, and the paper's claim is not that it is sophisticated but that it is reproducible. Nothing in Part VI or Appendix A.9 requires a figure that the procedures stated in this paper do not produce.

A.9 The recall link and jam budget

Earlier versions of this paper computed the recall receiver's *power* to three significant figures and left its *link* as an assertion, which several readers identified as the gravest omission in the specification. This appendix closes it. Nothing here is measured: it is arithmetic over standard published propagation and spread-spectrum relations, and every input is stated so that a reader may substitute their own and recompute.

A.9.1 The propagation model, which decides the answer

The governing choice is the propagation model, and it is where a calculation of this kind is usually flattered. Free-space loss between isotropic antennas is

$$L_{\text{fs}} = 20 \log_{10} d + 20 \log_{10} f - 147.55 \quad (d \text{ in m, } f \text{ in Hz})$$

but free space is the wrong law for a concealed antenna at ground level. Where a direct ray and a ground-reflected ray both reach the receiver, they arrive in near-antiphase and cancel, and beyond the first Fresnel breakpoint

$$d_b = \frac{4\pi h_t h_r}{\lambda}$$

the two-ray plane-earth model governs:

$$L_{\text{pe}} = 40 \log_{10} d - 20 \log_{10} h_t - 20 \log_{10} h_r$$

Two properties of that expression are counter-intuitive and both are load-bearing. **Loss rises at 40 dB per decade of range rather than 20**, so at 10 km the two models differ by about 50 dB—a factor of one hundred thousand in required power. And **frequency does not appear**, so moving band does not help; it moves only the breakpoint. What moves the answer is antenna height at each end, transmit power, and receiver sensitivity.




Figure 5. Free-space against two-ray plane-earth loss, for a station antenna at 10 m and a platform antenna at 0.5 m. Beyond the breakpoint the two diverge at 20 dB per decade, reaching about 50 dB at 10 km. Using the free-space law for a ground-level concealed antenna would understate the required power by that margin.

The received power follows, with L_{ex} a lump-sum allowance for foliage and terrain clutter and M a fade margin:

$$P_r = P_t + G_t + G_r - L_{\text{pe}} - L_{\text{ex}} - M$$

and the range at which a given receiver sensitivity S is met inverts to

$$d_{\text{max}} = 10^{B/40}, \quad B = P_t + G_t + G_r - S - L_{\text{ex}} - M + 20 \log_{10} h_t + 20 \log_{10} h_r$$

Inputs. $f = 150$ MHz; $h_t = 10$ m; $h_r = 0.3$ to 2 m; $G_t = 6$ dBi, $G_r = 0$ dBi; $L_{\text{ex}} = 12$ dB; $M = 10$ dB.

The model reproduces a fielded system’s published range, which is the only external check available. A fielded seismic unattended ground sensor node transmits at 15 mW from near ground level to a gateway on a raised antenna, and its manufacturer states a gateway-to-sensor standoff of one to five kilometres with options to ten.¹⁷⁶ Running the expression above over that geometry—the reverse of the recall link, and therefore the favourable one—returns about 0.7 to 1.9 km depending on the gateway’s receiver and antenna. The manufacturer’s stated lower bound falls inside that range and the higher options follow from a taller mast and a directional antenna, so the model is not obviously pessimistic.

A.9.2 The result, and the change it forces on the specification




Figure 6. Range achieved against transmit power at the control station, for a platform antenna at 0.5 m. The shaded band is the design case at \$15. The always-on wake-up stage does not reach it at any power a covert station would use; the gated receiver reaches it at twenty watts.

Station power	Always-on wake-up stage, −90 dBm	Gated receiver in a scheduled window, −120 dBm
5 W	1.3 km	7.5 km
20 W	1.9 km	10.6 km
100 W	2.8 km	15.8 km
500 W	4.2 km	23.7 km

[Derived. Platform antenna at 0.5 m; at 2 m the ranges roughly double, at 0.3 m they fall by about a quarter, since height enters squared at each end.]

¹⁷⁶Textron Systems, ‘MicroObserver UGS System’ (n 83), for the transmit power and standoff figures used as the model’s external check; and Agassounon and Steadman (n 84) for the 15 mW figure. The check compares a published range against a computed one and is not a measurement of either.

The design case at §15 does not close on the always-on stage. That case states a control station five to fifteen kilometres to the rear. On the wake-up receiver’s cited -90 dBm sensitivity, twenty watts at the station reaches under two kilometres, and five hundred watts—which is not a covert station—reaches four. The always-on stage cannot carry recall at the stated range and this paper should not have implied that it could.

It closes on the gated receiver, and that is a specification consequence rather than a rescue. A capable receiver live during a scheduled window, at -120 dBm, reaches ten kilometres on twenty watts and sixteen on a hundred. §15.9 already specifies that the platform “listens on a schedule that forms part of the cryptographic material”, and §17.3 budgets a receiver duty-cycled one second in sixty. What the link budget establishes is that **the schedule is not a power-saving convenience; it is the mechanism by which recall works at all.**

From which follows the result that matters outside engineering. **The latency of recall is set by the listening schedule, not by the link.** An operator cannot stop the system at an arbitrary instant at ten kilometres; they can stop it at the next window. On the one-second-in-sixty schedule at §17.3 the worst case is sixty seconds *to reach a window*—and §25.2 shows that reaching a window is not receiving a command. An authenticated message of two hundred bits at the rate its jam resistance requires takes a further two hundred seconds, so the worst case at the design point is **260 seconds**, and on an hourly window chosen to save power it is an hour and three minutes. Neither is instantaneous; a specification that states neither the window nor the message rate has not stated how timely its recall is; and §51 takes that to the treaty text.

A.9.3 Jamming

The recall message is short and, per §25, may be repeated. The information rate is therefore a design variable rather than a constraint, and processing gain follows it directly:

$$G_p = 10 \log_{10} \left(\frac{B}{R} \right)$$

Against a spread bandwidth $B = 1$ MHz, a rate R of one bit per second yields 60 dB. §25.2 fixes what that costs: an authenticated command is two hundred bits, so sixty decibels of gain is bought with two hundred seconds of reception, and eliminating R between the two relations gives $G_p = 10 \log_{10}(B/L) + 10 \log_{10} t_{\text{msg}}$ —**3 dB for every doubling of delivery time**. A jammer must overcome that gain, and the standoff at which it can be obtained by setting the jam-to-signal ratio at the receiver equal to G_p less the required margin.

Jammer power	1 bit/s, $G_p = 60$ dB	100 bit/s, $G_p = 40$ dB
10 W	0.33 km	1.06 km
100 W	0.59 km	1.88 km
1 kW	1.06 km	3.34 km

[Derived, against the gated receiver at 0.5 m, station at 20 W, 10 dB required margin.]

To deny recall, a jammer must be placed within roughly a kilometre of an emplacement whose location it does not know, and must be left there. That sharpens §27.1 rather than contradicting it: jamming a command-to-stop channel does deny recall while leaving function untouched, but the jammer must be local, persistent and itself emitting, which is an expensive and locatable commitment on ground the system exists to deny. Against the permission-to-continue




Figure 7. The standoff within which a jammer must sit to deny a five-kilometre recall link, against jammer power, at two information rates. Lowering the rate—which the message length permits—moves the whole curve down.

channel of §25.3 the same commitment buys the jammer a platform that safes itself, which is a poor return for a locatable emitter left on denied ground. An adversary close enough to jam the recall channel is close enough to be within the denied ground, and §27.2 records that clearance is the counter-measure that works.

A.9.4 What this appendix does not establish

It is a desk calculation over a flat-earth model with a single lump-sum allowance for clutter. It says nothing about diffraction over specific terrain, about urban multipath, about a jammer using a directional antenna or an airborne platform, or about the probability that any particular message is received. The experiment at A.2 stands unchanged and is now better specified: measure command delivery probability, at the geometry above, against a stated jammer, with the listening schedule as the independent variable.

A.10 Supplementary material and the reproducibility package

Every derivation in Part VI, §25 and Appendices A.8 and A.9 is reproducible from the expressions stated in the text, and each is simple enough to be checked by hand or in a spreadsheet. **The calculation is also supplied as code, because a two-hundred-thousand-draw sample cannot be reproduced from prose and because a reader is entitled to substitute an input and see what moves.**

The archive. The reproducibility archive of v2.12.2, published with this paper as the-emplaced-machine-reproducibility-v2.12.2.zip (in the deposit record, Reproducibility - The Emplaced Machine - v2.12.2.zip), remains controlling because none of v2.12.3, v2.12.4 and v2.12.5 changes any model input or code: v2.12.3 only corrected the displayed rounding of a value the archive computes as 1.216670 Wh/day, v2.12.4 changed only the treaty citations and the prose that depends on them, and

v2.12.5 corrects only prose, notes and cross-references. Its SHA-256, 464974348daded7c61a4b226ea009bc93aed390f4b4818ade327db0a46c56f3f, is also stated in the deposit record’s description, at <https://doi.org/10.5281/zenodo.23073940>; the archive contains MANIFEST.sha256, listing the digest of every file including the generated figures, verifiable with `sha256sum -c MANIFEST.sha256`. The reproducibility archive (v2.12.2) remains under CC BY 4.0, as its README states.

File	What it produces
<code>parameters.py</code>	the single canonical parameter file ; every value with its class
<code>budget_and_cascade.py</code>	§17.3–§17.4: the budget, the elasticities, the fixed-seed Monte Carlo
<code>detection_envelope.py</code>	§17.5: the constrained range, with the two illustrative sensing relations isolated at the top of the file
<code>modality_variants.py</code>	§15.5, §17.3 and §17.6: E_m in all four architectural forms, the sector-sweep cost of §15.5.1.3 over the exact rectangular solid angle, the public proxy envelope, the tier-one robustness table of §17.3, the amortised survey cost, the weather transfer, R_E , η , Φ , the common gate and Δ_m . F1_modality raises rather than defaulting when q_{survive} is absent.
<code>recall_coupling.py</code>	§25.2: the timeliness and jam-resistance coupling
<code>propagation_and_jamming.py</code>	Appendix A.9: the link and jam budget
<code>make_figures.py</code>	regenerates all seven manuscript figures and the four brief figures
README.md	environment, dependency versions, commands, units, expected outputs
MANIFEST.sha256	digests of code, inputs and generated figures

One canonical parameter file. No value is repeated by hand across scripts, and the manuscript’s tables are checked against it rather than transcribed from it. `python3 parameters.py` prints every input with its class: *assumption* where the author stipulated it, *component* where it is a published figure for a class of part, *design* where it is a design choice, *stipulation* where it is adopted from an external instrument. **No value in the package is a measurement made by this paper**, and the package says so in the same place it says the values.

The sampling, stated to the standard a reader needs to reproduce it. Seed 20260826; generator `random.Random(seed).triangular` from the CPython standard library; 200,000 draws; triangular on each variable with its mode at the design point; bounds as listed in `parameters.RANGES`; sampled independently; quantiles by `statistics.quantiles(data, n=100, method='exclusive')`, the fifth and ninety-fifth percentiles being elements 4 and 94. **The bounds are judgement bounds and not hard bounds**—nothing measured establishes that a value outside them is impossible—and **the independence assumption is questionable in one specific way**, since a noisy site plausibly degrades F_1 and q_2 together and a reader who models that correlation obtains a wider upper tail than §17.4 reports.

And a rule about disagreement, which replaces a worse one. An earlier version said that where the supplement and the text disagree, the supplement is right. That was wrong: neither silently governs the other. **A discrepancy between code, equation, table or prose is an**

error requiring correction before release; all four must be generated from, or checked against, the same declared inputs.

Every table and figure caption in this paper identifies whether its values are measured, sourced, stipulated or derived. Figure 3 is marked **ILLUSTRATIVE MODEL — NOT FIELD DATA** because it is a plot of stipulated sensing relations, and A.2 and A.11 are the experiments that would replace them with measured ones.

A.11 Active-sensor envelope and modality-substitution test

§15.5 makes two claims that can be tested without a weapon: that the usable envelope differs materially by sensing modality, and that a common downstream gate prevents a stronger modality from enlarging authority. **This is the highest-value experiment in the appendix**, because the second claim can be tested on an unarmed rig, requires no long-range or controlled sensor, and settles a safety property rather than a performance one.

Build: an instrumented static rig carrying passive electro-optical and thermal sensing, an available 3D lidar or a suitable recorded point-cloud source, a common time and pose reference across all sensors, energy measurement at each sensor and compute rail, and **a dummy effect-request sink placed behind the common authority gate**. No weapon, no release mechanism, no effector.

Modes. Test at least six:

#	Mode	What it isolates
1	passive EO/thermal only	the baseline envelope
2	one-time lidar survey plus passive operation	whether survey improves C_{pose} without changing C_{cue}
3	event-driven lidar confirmation	the priced confirmatory branch
4	fused passive and lidar classification	whether fusion changes $P_{\text{ID} D}$ or only P_D
5	automatic transition between passive and lidar paths	whether the transition itself is safe
6	lidar path presented without the camera path	the substitution case §15.5.3 describes

Measure. For each modality m , configuration θ_m , range and environment z : P_D ; $P_{\text{ID}|D}$; the relevant false-positive and false-release surrogate rates; F_1 and q_2 where the cascade applies; energy per event and per day under the declared duty cycle; latency; the emission-signature vector Σ_m ; pose-validity performance; $C_{\text{cue},m}$, $C_{\text{ident},m}$ and the resulting U_m ; and whether an automatic mode transition ever produces an effect request outside A_H .

Measure the invocation energy, term by term, because the public proxy calculation at §17.3 cannot establish any of it. For every tested mode, and separately for each modality m :

Quantity	Symbol	Why it cannot be inferred
Cold-start energy per invocation	$E_{\text{start},m}^{\text{cold}}$	no source examined reports it differs from cold start by the thermal state, which the duty cycle sets
Warm-start energy per invocation	$E_{\text{start},m}^{\text{warm}}$	

Quantity	Symbol	Why it cannot be inferred
Time to usable data after power-on	$t_{\text{ready},m}$	sets the floor on confirmatory latency, which enters §16.3's readiness metric
Stabilisation or calibration energy	included in $E_{\text{start},m}$	automated bias control against ambient temperature is documented ¹⁷⁷ and not costed
Electrical power while observing	$P_m(z, \text{activity})$	measure against scene activity and weather, not as a constant ; §15.5.1.2
Observation-window duration	$t_{\text{obs},m}(z)$	weather-dependent through retries
Modality-specific processing energy	$E_{\text{proc},m}$	point-cloud processing is a separate rail from the sensor
Retry count and retry energy	$n_{\text{retry}}(z), t_{\text{retry}}$	the mechanism by which weather enters a confirmatory architecture
Shutdown energy	$E_{\text{stop},m}$	folded into nothing at present
Invocation fraction under the declared policy	α_m	a property of the cascade, not of the sensor
Pointings, dwell and step-settle time for one sector sweep	$N_{\text{point}}, t_{\text{point}}$	§15.5.1.3 computes a floor from a published field of view; measure the real one
Receiver dissipation against array activity	P_m vs events/frame	the coupling at §15.5.1.2, which no source quantifies for a fielded configuration

Measure the weather transfer, and report the stages separately. The single most important measurement in this appendix is the one that converts a raw-return density into a system-level cue rate, because the paper refuses to assume it and the supplement refuses to compute without it:

Quantity	Symbol	Note
Raw unwanted-return density by condition	$\lambda_m(x, z)$	detections per second per square metre
Area and range interval over which it was measured	$A_m, [r_0, r_1]$	report both, always ; the published automotive figure is a near-field one ¹⁷⁸
Fraction removed at each filtering, clustering and tracking stage	—	stage by stage, not as a single number
Filter and tracking survival probability	$q_{\text{survive},m}(z, \theta)$	the transfer coefficient; the appendix exists for this
System-level false-cue rate where the modality may cue	$F_{1,m}(z)$	only meaningful for the cueing architecture
False confirmation, false rejection and abstention rates	—	the confirmatory architecture's equivalent, and a different quantity
Weather-dependent dwell and retries	$t_{\text{obs},m}(z), n_{\text{retry}}(z)$	carried back into E_m

¹⁷⁷Entwistle and others (n 39), on automated bias control against ambient temperature change.

¹⁷⁸Rasshofer, Spies and Spies (n 45), for the near-field measurement geometry that any replication must state.

Quantity	Symbol	Note
Detection and conditional identification by range and weather	$P_D, P_{ID D}$	by range bin, in each z
Uncertainty intervals and trial counts	—	report both or report nothing ; §17.4’s practice applies here

And two comparisons that decide claims this paper makes and cannot presently support. Where the hardware permits, run the same rig at a silicon-compatible wavelength and at an eye-safe one, and report the dark count rate, the resulting F_1 and the resulting N_3 at each: §15.5.1.4 asserts a direction and no magnitude, and this is the measurement that supplies one. And run the foliage case from a single fixed viewpoint against the same scene observed from several, to measure what §15.5.1.1 asserts qualitatively—that the occluded volume is irreducible from one position, and that the near-field desensitisation which recovers ground returns costs detection probability on the near object.

Required negative tests. These are the ones that decide the safety claim, and each must be run to failure rather than to success:

Negative test	Required behaviour
Candidates outside A_H detectable only by the stronger modality	the gate blocks every effect request
Pose-validity input removed or made stale	the gate refuses the affected function
Lidar classifier produces a high-confidence track the camera never saw	authority does not expand
Modality switched automatically after τ_H	the interval origin remains τ_H
One modality denied or corrupted	the system contracts, refuses or safes per the declared policy, and does not fall through to an unvalidated branch

Decides. Two claims, and the second is the stronger.

First, whether the energy- and emission-constrained usable envelope differs materially by modality. If the measured Δ_m of §15.5.5 sits inside the validated bounds, the substitution is a component change and the paper’s materiality test is uninteresting. If it does not—and the adverse-weather literature at §15.5.1 suggests it will not, in either direction—then a sensor change is a system change and must be reviewed as one.

Second, whether the common authority gate holds. This is the easier and the more valuable of the two, because it needs no access to a classified or long-range sensor: **a high-performance military lidar can be represented by a measured parameter set fed to the same architecture, without changing either the design or the proof obligation.** If a rig with a dummy effect output can be made to issue a request outside A_H through any modality path, the gate is not where the specification says it is, and §15.5.4 is a proposal rather than a property.

Reference List

This consolidated list supplements, but does not replace, the numbered footnotes. Evidentiary qualifications, access status, pinpoint references and source-specific caveats remain in the notes and should be read with the corresponding entry.

Technical and scientific literature

A H Lang and others, ‘PointPillars: Fast Encoders for Object Detection From Point Clouds’ in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition 12689–12697 doi:10.1109/CVPR.2019.01298

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

B F Aull, E K Duerr, J P Frechette, K A McIntosh, D R Schuette and R D Younger, ‘Large-format Geiger-mode avalanche photodiode arrays and readout circuits’ (2018) 24(2) IEEE Journal of Selected Topics in Quantum Electronics, article 3800510 doi:10.1109/JSTQE.2017.2736440

B F Aull, ‘Geiger-mode avalanche photodiode arrays integrated to all-digital CMOS circuits’ (2016) 16(4) Sensors 495 doi:10.3390/s16040495

‘Convergence and Emission Control’ Military Review (November–December 2023) 59 [online] accessed 5 August 2026

D Fried, T Mangum and G Stamnes, ‘2024 Annual GeoChallenge: showcasing innovation with Geiger-mode lidar’ (2025) 91(9) Photogrammetric Engineering & Remote Sensing 535–539.

D Seto, B H Krogh, L Sha and A Chutinan, ‘The Simplex architecture for safe online control system upgrades’ in Proceedings of the 1998 American Control Conference 3504–3508 doi:10.1109/ACC.1998.703255

David H Goldberg, Andreas G Andreou, Pedro Julián, Philippe O Pouliquen, Laurence Riddle and Rich Rosasco, ‘VLSI implementation of an energy-aware wake-up detector for an acoustic surveillance sensor network’ (2006) 2(4) ACM Transactions on Sensor Networks 594–611 doi:10.1145/1218556.1218562

Deng Yang and Jiahao Zhao, ‘Acoustic wake-up technology for microsystems: a review’ (2023) 14(1) Micromachines 129 doi:10.3390/mi14010129

Don J Torrieri, Principles of Spread-Spectrum Communication Systems (Springer).

George Succi, Gervasio Prado, Robert Gampert, Torstein Pedersen and Hardave Dhaliwal, ‘Problems in seismic detection and tracking’ in Unattended Ground Sensor Technologies and Applications II (SPIE 2000).

Gökhan 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.

Harry Urkowitz, ‘Energy Detection of Unknown Deterministic Signals’ (1967) 55 Proceedings of the IEEE 523.

IEC 60825-1, Safety of laser products—Part 1: Equipment classification and requirements.

J Lawrence, D B Kelley and D G Fried, ‘Ground estimation under foliage with Geiger-mode lidar’ in Laser Radar Technology and Applications XXX (2025) Proc SPIE doi:10.1117/12.3053936

Jorge Rivero, ‘Networked for War: Lessons from Ukraine’s Ground Robots’ (Modern War Institute, 9 March 2026) [online] accessed 5 August 2026

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

‘Leveraging Energy Harvesting and Wake-Up Receivers for Long-Term Wireless Sensor Networks’ (2018) 18 *Sensors* 1578.

M A Albota and others, ‘Three-dimensional imaging laser radars with Geiger-mode avalanche photodiode arrays’ (2002) 13(2) *Lincoln Laboratory Journal* 351–370.

M A Itzler and others, ‘Geiger-mode APD single-photon cameras for 3D laser radar imaging’ in 2014 *IEEE Aerospace Conference* 1–12 doi:10.1109/AERO.2014.6836476

M Bijelic, T Gruber, F Mannan, F Kraus, W Ritter, K Dietmayer and F Heide, ‘Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather’ in 2020 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* 11679–11689 doi:10.1109/CVPR42600.2020.01170

M Entwistle, M A Itzler, J Chen, M Owens, K Patel, X Jiang, K Slomkowski and S Rangwala (Princeton Lightwave Inc), ‘Geiger-mode APD camera system for single-photon 3-D LADAR imaging’ in *Advanced Photon Counting Techniques VI* (2012) *Proc SPIE* 8375, 83750D doi:10.1117/12.921004

Priyankar Choudhary, Neeraj Goel and Mukesh Saini, ‘A Survey on Seismic Sensor based Target Detection, Localization, Identification, and Activity Recognition’ (2023) *ACM Computing Surveys*, doi:10.1145/3568671.

R H Rasshofer, M Spies and H Spies, ‘Influences of weather phenomena on automotive laser radar systems’ (2011) 9 *Advances in Radio Science* 49–60 doi:10.5194/ars-9-49-2011

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

S Cova, M Ghioni, A Lacaita, C Samori and F Zappa, ‘Avalanche photodiodes and quenching circuits for single-photon detection’ (1996) 35(12) *Applied Optics* 1956–1976 doi:10.1364/AO.35.001956

S E Johnson, ‘Foliage penetration optimization for Geiger-mode avalanche photodiode lidar’ in *Laser Radar Technology and Applications XVIII* (2013) *Proc SPIE* 8731, 87310C doi:10.1117/12.2018911

Y Zhou and O Tuzel, ‘VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection’ in 2018 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* 4490–4499 doi:10.1109/CVPR.2018.00472

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.

Legal and policy materials

Christopher I. V. Farmer, ‘A Robot Left in a Field Is a Mine the Law Cannot See’ doi:10.5281/zenodo.21971256

Christopher I. V. Farmer, ‘The Interval, Not the Autonomy’ doi:10.5281/zenodo.21970778

Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons, Amended Protocol II (3 May 1996), in force 3 December 1998.

Convention on the Prohibition of the Use, Stockpiling, Production and Transfer of Anti-Personnel Mines and on their Destruction (Ottawa, 18 September 1997).

Final document of the Twenty-Seventh Annual Conference of the High Contracting Parties to Amended Protocol II, Geneva, 11 November 2025, UN Doc CCW/AP.II/CONF.27/5.

‘General Purpose Robots Should Not Be Weaponized: An Open Letter to the Robotics Industry and our Communities’ (6 October 2022) [online] accessed 26 August 2026

Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, chair’s summary, UN Doc CCW/GGE.1/2026/WP.2 (1 April 2026).

Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, Report of the 2024-2025-2026 sessions, UN Doc CCW/GGE.1/2026/CRP.1 (7 July 2026) para 36(d) [online] accessed 26 August 2026

Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, Report of the 2024-2025-2026 sessions, UN Doc CCW/GGE.1/2026/CRP.1/Rev.1 (3 September 2026) [online] accessed 1 October 2026

Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, Report of the 2024-2025-2026 sessions, UN Doc CCW/GGE.1/2026/3 (advance copy, 14 September 2026) [online] accessed 15 October 2026

Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, rolling text, status 18 December 2025, section III para 7(D) [online] accessed 26 August 2026

Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, rolling text, status 5 June 2026, para 12(D).

Hitoshi Nasu, ‘The Kargu-2 Autonomous Attack Drone: Legal and Ethical Dimensions’ (Lieber Institute, 10 June 2021) [online] accessed 26 August 2026

Human Rights Watch, ‘Background Briefing on Landmine Use in Ukraine’ (15 June 2022) [online] accessed 26 August 2026

International Committee of the Red Cross, A Guide to the Legal Review of New Weapons, Means and Methods of Warfare (ICRC 2006).

International Committee of the Red Cross, ICRC Position on Autonomous Weapon Systems (12 May 2021) [online] accessed 26 August 2026

Letter dated 8 March 2021 from the Panel of Experts on Libya established pursuant to resolution 1973 (2011) addressed to the President of the Security Council, UN Doc S/2021/229 (8 March 2021) para 63 [online] accessed 26 August 2026

Lieber Institute, ‘Terrain Denial Fires on the Modern Battlefield’ [online] accessed 5 August 2026

Meeting of the High Contracting Parties to the Convention on Certain Conventional Weapons, final document, UN Doc CCW/MSP/2025/8.

Paul Scharre and Michael C Horowitz, *An Introduction to Autonomy in Weapon Systems* (Center for a New American Security, February 2015).

Paul Scharre, 'Autonomy, "Killer Robots," and Human Control in the Use of Force' (Just Security, 9 July 2014) [online] accessed 26 August 2026

Protocol on Blinding Laser Weapons (Protocol IV to the Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons), Vienna, 13 October 1995.

'Renewed call by the UN Secretary-General and ICRC President to adopt rules on autonomous weapons' (International Committee of the Red Cross, 25 August 2026) [online] accessed 26 August 2026

United Nations Treaty Collection, depositary notification C.N.421.2025.TREATIES-XXVI.5 [online] accessed 26 August 2026

United States statement at the Group of Governmental Experts on Lethal Autonomous Weapons Systems, agenda item 5(d), 5 August 2021 [online] accessed 26 August 2026

United States statement to the Group of Governmental Experts on Lethal Autonomous Weapons Systems, 'Consideration of the Human Element in the Use of Lethal Force' (27 March 2019) [online] accessed 26 August 2026

US Department of Defense Directive 3000.09, *Autonomy in Weapon Systems* (25 January 2023) para 1.2.d(2) [online] accessed 26 August 2026

Vincent Boulanin and Maaïke Verbruggen, *Mapping the Development of Autonomy in Weapon Systems* (Stockholm International Peace Research Institute 2017) [online] accessed 26 August 2026

Zachary Kallenborn, 'If a killer robot were used, would we know?' (Bulletin of the Atomic Scientists, 4 June 2021) [online] accessed 26 August 2026

Government and programme records

Center for Army Lessons Learned, US Army Combined Arms Center, No 26-1140 (March 2026).

Department of National Defence (Canada), Ground Surveillance Radar, B-GL-317-015/FP-001 (CFP317(15)) [online] accessed 26 August 2026

Director, Operational Test and Evaluation, FY2001 annual report, 'Wide Area Munition' [online] accessed 26 August 2026

Director, Operational Test and Evaluation, FY2010 annual report, Spider XM7 Network Command Munition [online] accessed 26 August 2026

Eric D Haider, *Unattended Ground Sensors and Precision Engagement*, Master's thesis, Naval Postgraduate School, Monterey (December 1998), DTIC AD-A359912.

Federation of American Scientists, 'M131 Modular Pack Mine System (MOPMS)' [online] accessed 26 August 2026

Federation of American Scientists, 'M93 HORNET' [online] accessed 26 August 2026

Federation of American Scientists, 'Remote Battlefield Sensor System (REMBASS) / Improved Remote Battlefield Sensor System (I-REMBASS)' [online] accessed 6 August 2026

John K Hawley, *Patriot Wars: Automation and the Patriot Air and Missile Defense System* (Center for a New American Security, January 2017).

Lawrence R Edwards, Sterling S Pilette, Billy E Biggs and Harold Martinek, *The Effect of Workload on Performance of Operators Monitoring Unattended Ground Sensors*, Technical Paper 321, US Army Research Institute for the Behavioral and Social Sciences (September 1978), DTIC AD-A061694.

Letter from Senator Elizabeth Warren and Representative Donald Norcross to the Secretary of Defense, 12 August 2026 [online] accessed 26 August 2026

Letter from Senator Elizabeth Warren to the Secretary of Defense, 1 February 2022 [online] accessed 26 August 2026

National Research Council, *Naval Mine Warfare: Operational and Technical Challenges for Naval Forces* (National Academies Press 2001).

National Security and Defence Council of Ukraine, 'Results of Ukraine's Defense Industry in 2025: FPV Drones' [online] accessed 5 August 2026

NATO Science and Technology Organization, *Disposable Multi-Sensor Unattended Ground Sensor Systems for Detecting Personnel*, STO Technical Report TR-SET-158, AC/323(SET-158)TP/578, DTIC AD-A626493.

R M Marino, T Stephens, R E Hatch, J L McLaughlin, J G Mooney, M E O'Brien, G S Rowe, J S Adams, L Skelly, R C Knowlton, S E Forman and W R Davis, 'A compact 3D imaging laser radar system using Geiger-mode APD arrays: system and measurements' in *Laser Radar Technology and Applications VIII* (2003) Proc SPIE 5086 doi:10.1117/12.501581

Sterling Pilette, Bill Biggs and Harold Martinek, *The Value of Special Training and Job Aids for Improving Unattended Ground Sensor Operator Performance*, Technical Paper 304, US Army Research Institute for the Behavioral and Social Sciences (August 1978), DTIC AD-A061046.

US Army and US Marine Corps, ATP 3-90.8 / MCWP 3-17.5: *Combined Arms Countermobility Operations* (September 2014).

US Army, FM 20-32: *Mine/Countermining Operations*, ch 3 [online] accessed 26 August 2026

US Army, FM 5-102: *Countermobility*, ch 7 ('Denial Operations') [online] accessed 5 August 2026

US Army, FM 90-7: *Combined Arms Obstacle Integration*.

US Army, 'Adapting to Multi-Domain Battlefield: Developing Emissions Control SOP' [online] accessed 5 August 2026

US Army, 'Along came a spider' (9 October 2012) [online] accessed 26 August 2026

US Army, 'Critical events in developing the concept of operations for terrain-shaping obstacles' (2019) [online] accessed 26 August 2026

US Army, 'Soldiers shape the battlefield with next generation of terrain-shaping munitions' (army.mil, 2019) [online] accessed 26 August 2026

US Department of Defense, Office of the Inspector General, *Acquisition of the Wide Area Munition*, Report No D-2002-011 (24 October 2001) [online] accessed 26 August 2026

US Department of Defense, Performance Metrics for Low Probability of Intercept Communication Systems (Defense Technical Information Center, ADA277346) [online] accessed 5 August 2026

US Government Accountability Office, Weapon Systems Testing: Reorganization of DOD's Office of the Director, Operational Test and Evaluation, GAO-26-108859 (30 June 2026) [online] accessed 26 August 2026

US Naval Institute, 'Mission Command Means Emissions Control' (Proceedings, May 2021) vol 147/5/1,419 [online] accessed 5 August 2026

US Navy, 'Aegis Weapon System' fact file [online] accessed 26 August 2026

US Navy, 'MK 15 Phalanx Close-In Weapon System (CIWS)' fact file [online] accessed 26 August 2026

William Agassounon and Robert Steadman, The Application of Unattended Ground Sensors to Human Socio-Cultural Behavior Monitoring (17 October 2012), DTIC AD-A581145.

Manufacturer, industry and press sources

Anduril, 'Anduril Launches Extended Range Sentry Tower (XRST) to Advance Counter Intrusion Missions' [online] accessed 5 August 2026

Anduril, 'Sentry' [online] accessed 5 August 2026

Army Recognition, 'Chinese Army Demonstrates Machine Gun-Equipped Robotic Dogs in Joint Exercise with Cambodia' [online] accessed 5 August 2026

ASD News, 'Anduril Deploys 300th Autonomous Surveillance Tower (AST), Advancing Capability Border Security' [online] accessed 5 August 2026

BAE Systems, 'Geiger-mode lidar camera' product page [online] accessed 27 August 2026

Belfer Center for Science and International Affairs, 'Starlink and the Russia-Ukraine War: A Case of Commercial Technology and Public Purpose?' [online] accessed 5 August 2026

D Fried and K Reichel-Vischi (3DEO Inc), 'Next-generation Geiger-mode lidar systems' LiDAR Magazine (30 December 2024) [online] accessed 27 August 2026

David Hambling, 'Ukraine Turns To AI-Controlled Guns To Stop Russian Shahed Drones' (Forbes, 27 May 2025) [online] accessed 5 August 2026

Defence Blog, 'U.S. Marine Corps taps an AI-powered gun against cheap drones' (21 July 2026) [online] accessed 24 August 2026

Delta-Q Technologies, 'Building a Lights-Out Warehouse with Autonomous Charging Systems' [online] accessed 5 August 2026

Edge AI Stack, 'Jetson Orin Nano Power Consumption: Idle, 15W, MAXN Benchmarks' (2026) [online] accessed 5 August 2026

Erico Guizzo and Evan Ackerman, 'Do We Want Robot Warriors to Decide Who Lives or Dies?' (IEEE Spectrum, 31 May 2016) [online] accessed 26 August 2026

Euromaidan Press, 'A Russian drone chose its own target and killed three. The code that told it what to look for was readable' (25 August 2026) [online]

Evan Ackerman, ‘Q&A: Ghost Robotics CEO on Armed Robots for the U.S. Military’ (IEEE Spectrum, 18 October 2021) [online] accessed 24 August 2026

Forecast International, ‘M93 Hornet Wide Area Munition’ (archived programme report) [online] accessed 26 August 2026

GlobalSecurity, ‘AEGIS Core Processes’ [online] accessed 26 August 2026

GlobalSecurity, ‘MK 60 Encapsulated Torpedo (CAPTOR)’ [online] accessed 26 August 2026

GlobalSecurity, ‘Over 9,000 frontline missions in March: the Defence Forces continue to expand the use of unmanned ground vehicles’ [online] accessed 5 August 2026

GlobalSecurity, ‘Remote Battlefield Sensor System (REMBASS)’ [online] accessed 6 August 2026

GlobalSecurity, ‘Unattended Ground Sensors (UGS)’ [online] accessed 5 August 2026

Hesai Technology, ‘OT128 product specifications’ [online] accessed 27 August 2026

IEEE Spectrum, ‘A Robotic Sentry For Korea’s Demilitarized Zone’ [online] accessed 6 August 2026

Joseph Trevithick, ‘The U.S. Is Getting Back Into Naval Mine Warfare’ (The War Zone, 2019) [online] accessed 26 August 2026

Katie Livingstone, ‘Ukraine to field 25,000 ground robots in push to replace soldiers for frontline logistics’ Defense News (24 April 2026) [online] accessed 16 October 2026

Kharon, ‘At Unitree Robotics, a Star Chinese Firm, the Military Connections Keep Mounting’ [online] accessed 5 August 2026

LightPath Technologies, ‘Uncooled vs Cooled LWIR Cameras: A Guide’ [online] accessed 5 August 2026

Livox, ‘Mid-360 specifications’ [online] accessed 27 August 2026

Meduza, ‘NYT: A fully autonomous Russian AI drone run by an Nvidia minicomputer killed 3 civilians in the Ukrainian city of Zaporizhzhia’ (24 August 2026) [online]

Militarnyi, ‘Sky Sentinel: Ukraine to produce anti-aircraft turrets for intercepting drones and cruise missiles’ [online] accessed 5 August 2026

Militarnyi, ‘UGV With a Machine Gun Replaced Infantry and Held a Position for a Month and Half’ (Militarnyi, 15 November 2025) [online] accessed 5 August 2026

N R Jenzen-Jones and others, ‘Russian PTKM-1R top-attack anti-vehicle mine documented in Ukraine’ (Armament Research Services, 2022) [online] accessed 26 August 2026

Ouster, OS2 hardware user manual for Rev7 sensors (14 October 2024) [online] accessed 27 August 2026

PBS NewsHour, ‘Elon Musk’s refusal to provide Starlink support for Ukraine attack in Crimea raises questions for Pentagon’ (PBS NewsHour) [online] accessed 5 August 2026

Phihong, ‘Swappable Battery vs. Autonomous Docking: What Leading OEMs Are Choosing for Next-Gen Mobile Robots’ [online] accessed 5 August 2026

Pkcell, ‘Comprehensive Guide to Lithium Thionyl Chloride Batteries: Operating Principles, Passivation, and IoT Hybrid Power Solution’ [online] accessed 5 August 2026

PVKnowHow, ‘Ukraine Solar Panel Manufacturing: Market Insights Report’ [online] accessed 5 August 2026

Rosoboronexport, ‘PTKM-1R’ product page [online] accessed 26 August 2026

‘SGR-A1’ [online] accessed 6 August 2026

Simon Parkin, ‘Killer robots: The soldiers that never sleep’ (BBC Future, 16 July 2015) [online] accessed 26 August 2026

‘Starlink satellite services in Ukraine’ [online] accessed 5 August 2026

Stavroula Pabst, ‘Why would the Pentagon want to hide weapons testing results?’ (Responsible Statecraft, 20 August 2026) [online] accessed 26 August 2026

Tadiran Batteries, ‘Long-Life Lithium Batteries’ [online] accessed 5 August 2026

Textron Systems, ‘MicroObserver UGS System’ (product literature).

The Robot Report, ‘Army readies charging port for swarms of autonomous drones’ [online] accessed 5 August 2026

The War Zone, ‘Robot Dogs Now Have Assault Rifles Mounted On Their Backs’ (12 October 2021, updated 15 October 2021) [online] accessed 24 August 2026

ThingsLog, ‘Case study: DARPA N-ZERO near-zero-power battlefield sensors’ [online] accessed 6 August 2026

Throwflame, ‘Thermonator’, product page [online] accessed 24 August 2026

‘TM-83 mine’ (Wikipedia) [online] accessed 26 August 2026

Tom’s Hardware, ‘Nvidia Jetson Orin-guided Russian AI drone killed three civilians in Ukraine, forensic teams say’ (2026) [online]

‘Ukraine to buy 4.5 million FPV drones in 2025’ [online] accessed 5 August 2026

Ukrainian News, ‘Russia sharply increases drone production: plans to produce more than 7 million FPVs in 2026’ [online] accessed 5 August 2026

‘Unattended ground sensor’ (Wikipedia) [online] accessed 5 August 2026

Unitree Robotics, ‘Humanoid robot G1’, product page [online] accessed 24 August 2026

Unmanned Systems UK, ‘UGV Power Systems and Energy Management’ (15 June 2025) [online] accessed 5 August 2026

Datasets, code and supplementary materials

A Geiger, P Lenz and R Urtasun, ‘Are we ready for autonomous driving? The KITTI vision benchmark suite’ in 2012 IEEE Conference on Computer Vision and Pattern Recognition 3354–3361 doi:10.1109/CVPR.2012.6248074

H Caesar and others, ‘nuScenes: A Multimodal Dataset for Autonomous Driving’ in 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition 11618–11628 doi:10.1109/CVPR42600.2020.01164

Open Robotics, ROS 2 documentation [online] accessed 24 August 2026

Unitree Robotics, public repositories at [online] accessed 24 August 2026