MODULAR ANTI-DRONE CAGE ARMOUR ON MAIN BATTLE TANKS: AN ENGINEERING ASSESSMENT OF EFFECTIVENESS AGAINST FPV-DELIVERED SHAPED-CHARGE WARHEADS IN TRENCH WARFARE CONDITIONS

Emre Aydin

Defence Industry Agency
Ankara, Republic of Turkey
E-mail: emre.aydin@ssb.gov.tr

Original Scientific Article

Received: 17.03.2026. Approved: 21.05.2026.

DOI: https://doi.org/10.65932/military-studies-2026-1-2 **UDC**: 623.438.3:623.562

Abstract: First-person view (FPV) drones carrying repurposed shaped-charge warheads have inverted the cost-of-attrition logic of mechanised warfare, and crews have responded with improvised roof grids, the so-called 'cope cages', whose engineering value remains contested. This article asks a narrow question: under what geometric conditions does a roof-mounted, modular cage measurably reduce the perforation probability of a top-attack FPV-delivered shaped charge, and when does it fail to do so? The study integrates the grid-interaction probability model developed for slat armour against the RPG-7 threat with the empirically anchored standoff–penetration response of a representative warhead and with established jet-formation and penetration mechanics. Its original contribution is the Top-Attack Cage Effectiveness (TACE) index, a dimensionless composite that multiplies the warhead-disruption probability of a grid by the residual-penetration margin of any jet that survives initiation at the cage standoff, evaluated against the roof protection threshold. The index makes explicit a counter-intuitive result: because a shaped charge reaches peak penetration at a finite optimal standoff, about 317 mm against 278 mm at contact for the modelled warhead, a cage that merely adds standoff without disrupting the warhead can raise penetration rather than lower it. Illustrative computation across grid pitches of 30–60 mm and roof gaps of 250–450 mm shows TACE values ranging from below 0.30 for poorly placed coarse grids to above 0.80 for fine, disruption-dominated configurations. The analysis concludes that cope cages are defensible as the outermost layer of a graduated counter-UAS survivability architecture rather than as standalone protection, and that their design must be driven by warhead-disruption geometry, not standoff alone. All numerical TACE values are model-derived and require experimental validation before fielding.

Keywords: *cage armour, FPV drones, shaped charge, standoff distance, top-attack protection, main battle tank, survivability.*

INTRODUCTION

The armoured vehicle entered the third decade of the twenty-first century as the presumed master of the close battle and is leaving the first years of that decade as one of its most exposed actors. The proximate cause is small and cheap. A first-person view (FPV) quadcopter costing a few hundred dollars, fitted with a salvaged anti-tank grenade, can disable a platform whose unit cost is measured in millions, and it does so from an axis — the roof — that armour designers have historically treated as a residual concern (Mittal & Goetz, 2025; Kweera, 2023). Quantitative battlefield analysis of the Russia–Ukraine war records cost-exchange ratios so lopsided that they reorder procurement priorities rather than merely complicate them, with drone and counter-drone measures now shaping the tempo of mechanised operations (Mittal & Goetz, 2025). The tactical setting that concerns this article is the most demanding one: static or semi-static trench warfare, where vehicles loiter within dense FPV engagement envelopes for hours, lack the speed to spoil a terminal dive, and present their thinnest surfaces, turret roof, engine deck, hatch rings, to an operator who can choose the impact point.

Crews did not wait for doctrine. From late 2022 onward, improvised welded grids appeared over turret roofs across the theatre, christened 'cope cages' by sceptics who doubted they did anything but inconvenience the maintainer (Teixeira da Silva, 2023). The scepticism is understandable. Cage, slat, and bar armours were conceived for a different geometry: the roughly horizontal RPG-7 round striking the flank, where a lattice of bars can crush the warhead body, sever the fuze, or otherwise interrupt the detonation chain before a coherent jet forms (Coghe, 2022). Transplanting that lattice to the roof against a near-vertical FPV dive imports the hardware without importing the physics that justified it. Whether the transplant works is an engineering question, not a rhetorical one, and it has not been answered with the specificity the problem deserves.

The shaped charge that defines the threat is unforgiving and well understood. Detonation collapses a metallic liner into a hypervelocity jet whose tip can exceed 8 km/s, and whose penetration into steel scales with jet length, density, and the standoff at which it is allowed to stretch before impact (Walters-type virtual-origin treatments are revisited by Elek et al., 2024, and Jia et al., 2021). Penetration is not monotone in standoff. A jet must travel far enough to elongate yet not so far that it particulates and drifts off axis, so penetration rises to a maximum at an optimal standoff of a few charge diameters and declines thereafter (Dehestani et al., 2019; Żochowski et al., 2021). This single fact, an internal maximum in the standoff–penetration curve, is the pivot on which the entire cope-cage question turns, and it is routinely elided in popular and even professional commentary that treats 'more standoff' as synonymous with 'more protection'.

The central research question is therefore precise. Under what geometric conditions does a roof-mounted modular cage measurably lower the perforation probability of a top-attack, FPV-delivered shaped charge against a main battle tank in trench-warfare conditions, and under what conditions does it fail, or backfire? Three hypotheses structure the inquiry. The first holds that the dominant protective mechanism of a roof cage against an FPV shaped charge is warhead and fuze disruption, not standoff increase, and that a design optimised only for standoff can increase residual penetration. The second holds that there exists a bounded, computable geometric window, defined jointly by grid pitch relative to warhead diameter and by roof gap relative to the warhead's optimal standoff, within which cage effectiveness is maximised. The third holds that, even at its geometric optimum, a passive cage is insufficient as standalone protection and yields its value only as the outer tier of a layered counter-UAS survivability architecture that also includes soft-kill, hard-kill, and reactive elements (Castrillo et al., 2022; Li et al., 2023; Moldtmann et al., 2024).

The original contribution of this article lies in a single analytical instrument built to answer that question: the Top-Attack Cage Effectiveness (TACE) index. The index couples two bodies of work that the literature has kept apart. From slat-armour research it takes the probability that a grid disrupts an incoming warhead as a function of pitch, bar width, and obliquity (Coghe, 2022); from terminal-ballistics research it takes the residual penetration of any jet that survives initiation, read off the standoff–penetration curve at the cage's effective standoff (Żochowski et al., 2021; Elek et al., 2024). No prior study, to my knowledge, has joined the grid-disruption probability to the standoff-dependent residual penetration to define an optimal roof-cage geometry specifically for the top-attack FPV case; the cage-armour canon addresses the horizontal RPG threat, while the standoff studies treat the munition in isolation. TACE fills that gap, and it does so in a form simple enough to guide design trade-offs in the field.

The argument proceeds as follows. A review situates the cope cage within the slat-armour, reactive-armour, active-protection, and counter-UAS literatures and isolates the gap the index addresses. The methodology defines TACE and its inputs. A results section reports illustrative index values across a realistic design space. Three analytical sections then interpret those results, first the standoff paradox, then the disruption mechanism on a roof-mounted grid, then the modular-geometry and mass trade-offs peculiar to trench warfare, before a fourth situates the cage within the layered survivability stack. A conclusion returns to the three hypotheses.

LITERATURE REVIEW AND METHODOLOGY

Literature Review

The terminal-ballistics foundation of the problem is mature and quantitatively rich. The collapse of a conical liner into a stretching jet, and the dependence of penetration on jet kinematics, are treated through virtual-origin and modified-Bernoulli models that recent work has refined and validated against experiment (Elek et al., 2024; Jia et al., 2021). The influence of liner material and geometry on jet coherence and penetration is documented across copper, zirconium, and layered designs, with zirconium and multi-layer liners offering distinctive trade-offs between density and reactivity (Elshenawy et al., 2022; Hao et al., 2023; Elshenawy et al., 2024). Numerical method itself has matured: smoothed-particle hydrodynamics now reproduces jet penetration into ceramic–metal stacks with fidelity sufficient for design exploration (Chen et al., 2023), and Johnson–Cook constitutive treatments capture the strain-rate sensitivity of liner and target alike (Shi et al., 2022). Experimental and numerical penetration campaigns now routinely cross-validate one another, and they show that subtle liner features, cavity filling and charge configuration, shift penetration by tens of millimetres (Du et al., 2022; Żochowski & Warchoł, 2023). For the present purpose the decisive strand is the dependence of penetration on standoff. Both dedicated parametric studies (Dehestani et al., 2019) and warhead-specific experimental campaigns (Żochowski et al., 2021) confirm the non-monotone curve: penetration climbs with standoff to a maximum and then falls, and the location of that maximum, measured in charge diameters, governs whether added standoff helps or hurts.

The standoff effect is not a marginal correction. In the experimentally validated PG-7VM campaign, penetration into high-hardness armour rose from roughly 278 mm at the nominal configuration to about 317 mm at a standoff of 360 mm before declining, a swing of some 14 percent driven by standoff alone (Żochowski et al., 2021). A dynamic-condition treatment shows that target motion and obliquity further modulate the effective standoff each jet element experiences, complicating any single-number characterisation (Jia et al., 2021). The rear-initiated and waveshaped configurations studied elsewhere demonstrate how sensitively jet quality responds to initiation geometry (Du et al., 2021; Pyka et al., 2020), which matters because a cage that initiates a warhead off-axis or off-nominal alters not only standoff but jet formation itself.

The slat-armour literature supplies the second pillar. The defeat mechanism of a bar grid against a shaped-charge round is not, as folklore holds, the dissipation of an already-formed jet; it is the prevention of proper jet formation through mechanical insult to the warhead before detonation, crushing the body between bars, shorting or shearing the piezoelectric fuze, or deforming the liner so the collapse is asymmetric (Coghe, 2022). Coghe's central result is a methodology that computes the efficiency of slat, bar, net, and inertial systems as the probability that the fuze fails to contact a grid element capable of triggering or disrupting it, expressed as a function of impact geometry; net systems achieve the best mass efficiency but degrade sharply above roughly 15 degrees of vertical impact angle (Coghe, 2022). That degradation is precisely the regime a top-attack FPV occupies, and it is the empirical hook for the present analysis.

Reactive and active protection define the comparison class. Explosive reactive armour disrupts a jet by driving plates laterally across its path, and recent work has both revisited the governing mass-flux model (Mayseless et al., 2019) and shown that data-driven and Bayesian-optimisation methods can design minimum-weight cassettes against combined kinetic-energy and shaped-charge threats (Becker et al., 2022; Moldtmann et al., 2024). The penetration of explosively formed projectiles against reactive armour quantifies the residual threat after a reactive interaction (Nie et al., 2019), and the formation of such projectiles is itself acutely sensitive to liner geometry and thickness gradients (Yang & Lin, 2021; Ma et al., 2022). Emerging reactive-liner charges that couple jet penetration with internal blast further complicate the protection calculus (Zhang et al., 2022). Active protection systems engage the munition before impact; ensemble-learning methods now predict the angular disturbance of a moving launcher to time countermeasure ejection (Li et al., 2023), and combat-effectiveness modelling assesses their kill probability against diverse threats (Park et al., 2024). The counter-UAS literature, finally, maps the soft-kill and hard-kill option space, radio-frequency and satellite-navigation denial, directed energy, capture, and cooperative defensive drones, and its limits against autonomous terminal guidance (Castrillo et al., 2022; Lee et al., 2022; Silva et al., 2023).

Two adjacent strands close the frame. Survivability and blast research, including the instrumentation of crew compartments with microelectromechanical accelerometers and the design of occupant-protection countermeasures, supplies the systems perspective in which any add-on must be weighed against mass, mobility, and crew burden (Kciuk et al., 2022; Pratomo et al., 2020). And the only peer-reviewed treatment to address the tank 'umbrella' directly frames the cope cage as a contested expedient whose physics deserve scrutiny rather than dismissal (Teixeira da Silva, 2023). The gap is now visible. The slat-armour model quantifies disruption for the horizontal RPG; the standoff studies quantify penetration for the munition alone; no one has joined the two for the top-attack FPV case to ask where, geometrically, a roof cage helps. That junction is the contribution this article offers.

Research Methodology

I adopt an analytical-engineering design rather than an experimental one, for a reason I should state plainly: I have no access to a proof range, and fabricating empirical penetration figures would be worse than useless. The study therefore synthesises published, peer-reviewed data into a transparent model whose every input is traceable, and it labels every output as model-derived pending validation. The approach has three layers. The first is a parametric standoff–penetration function calibrated to the experimentally validated PG-7VM response (Żochowski et al., 2021) and cross-checked against the parametric trends reported for stand-off and liner thickness (Dehestani et al., 2019) and the virtual-origin treatment of Elek et al. (2024). The second is a grid-disruption probability adapted from the slat-armour efficiency methodology (Coghe, 2022), re-parameterised for near-vertical obliquity. The third is the composite index that joins them.

The Top-Attack Cage Effectiveness index is defined for a cage configuration C against a threat T as the probability that the roof survives a single terminal engagement. Let P_d denote the probability that interaction with the grid disrupts the warhead so that no effective jet forms, a function of grid pitch p, bar width w, projected warhead diameter d, and impact obliquity θ. Let s_{eff} denote the effective standoff at which an undisrupted warhead initiates, equal to the roof gap g plus the warhead's intrinsic standoff. Let $P_{res}(s_{eff})$ be the residual penetration read from the standoff–penetration curve at s_{eff}, and let P_{crit} be the roof's protection threshold, the penetration required to perforate the roof plate and any backing. The index is $TACE = P_d + (1 - P_d) \cdot \Phi(P_{crit} - P_{res}(s_{eff}))$, where Φ is a bounded protection function returning one when residual penetration falls safely below the threshold, zero when it exceeds it, and a graded value across the transition. The first term credits outright disruption; the second credits the residual margin of jets that survive initiation. The index lies in the unit interval, and a value near one denotes near-certain protection in a single engagement.

The design space is bounded by what trench-warfare retrofits can actually carry and weld. Grid pitch ranges over 30, 40, 50, and 60 mm, values bracketing the projected diameter of the warheads most often improvised onto FPV airframes, from PG-7 derivatives of roughly 70–93 mm cone diameter to smaller custom charges. Roof gap ranges over 250, 350, and 450 mm, the practical envelope between hatch clearance and the height at which a cage becomes a mobility and signature liability. Impact obliquity is set near vertical, between 70 and 90 degrees from the horizontal roof plane, the regime in which net efficiency is known to collapse (Coghe, 2022). The roof protection threshold P_{crit} is treated parametrically around the thin-roof values typical of main battle tanks, a deliberately conservative anchor given that survivability studies emphasise how little margin the roof carries (Kciuk et al., 2022; Pratomo et al., 2020).

Three limitations are intrinsic to this design and I prefer to foreground them rather than bury them. First, the disruption probability borrows a model validated for the RPG-7 body and fuze, and an FPV-borne warhead may present a different fuze type and a different body stiffness, so P_d is indicative rather than exact. Second, the standoff curve is calibrated to one well-documented warhead; other charges will shift the optimal standoff and the magnitude of the penetration swing, though the qualitative non-monotonicity is general (Dehestani et al., 2019; Elek et al., 2024). Third, the model treats a single engagement, whereas trench warfare imposes repeated attacks and cumulative damage that a one-shot index cannot capture. These constraints set the boundary of the claims that follow.

RESEARCH RESULTS

The empirical and computational findings organise into three blocks corresponding to the three hypotheses: the shape of the standoff–penetration response that governs the standoff term; the behaviour of the disruption probability across the grid design space; and the composite TACE values that emerge from their combination. I present them here without interpretation, which follows in the analytical sections.

The standoff–penetration response, calibrated to the validated PG-7VM data, is unambiguously non-monotone. Penetration into high-hardness armour stands near 278 mm at the nominal short standoff, rises to a maximum of about 317 mm at a standoff of roughly 360 mm, and declines beyond that point (Żochowski et al., 2021). Expressed against the warhead's cone diameter, the maximum sits at a few charge diameters, consistent with the parametric optimum reported independently (Dehestani et al., 2019) and with the virtual-origin prediction that penetration tracks jet stretch until particulation intervenes (Elek et al., 2024). Table 1 records the calibration points used in the model.

Standoff (mm)	Penetration into high-hardness steel (mm)	Regime
~0 (contact)	≈ 278	rising branch
~180	≈ 300	rising branch
~360 (optimum)	≈ 317	maximum
~540	≈ 305	declining branch
~720	≈ 285	declining branch

Table 1. Calibration points of the standoff–penetration response used in the model, after the experimentally validated PG-7VM campaign (Żochowski et al., 2021) and the parametric trend of Dehestani et al. (2019).

The values at non-optimal standoffs are interpolated from the reported maximum and endpoints and are presented as model inputs, not as independent measurements.

The disruption probability behaves as the slat-armour methodology predicts, but in the unfavourable part of its range. Adapting the grid-coverage logic of Coghe (2022) to near-vertical obliquity, P_d falls as pitch widens relative to warhead diameter, because a coarser grid offers a smaller chance that the fuze or body strikes a bar with sufficient resistance. For a projected warhead diameter near 85 mm, a 30 mm pitch yields a high disruption probability on the order of 0.7–0.8, a 40 mm pitch on the order of 0.55–0.65, a 50 mm pitch around 0.4–0.5, and a 60 mm pitch below 0.4, with all values depressed by 10–20 percent relative to the horizontal case once obliquity exceeds about 70 degrees, consistent with the documented collapse of net efficiency above 15 degrees of vertical angle (Coghe, 2022). Table 2 summarises the modelled disruption probabilities.

Grid pitch (mm)	Pitch / diameter	P_d (near-vertical)
30	0.35	0.70 – 0.80
40	0.47	0.55 – 0.65
50	0.59	0.40 – 0.50
60	0.71	0.30 – 0.38

Table 2. Modelled warhead-disruption probability P_d of a roof grid versus grid pitch, for a projected warhead diameter of ≈ 85 mm at near-vertical obliquity, adapted from the slat-armour efficiency methodology of Coghe (2022). Values are model-derived.

The composite index exposes the interaction the two preceding tables only imply. When the roof gap places an undisrupted warhead near its optimal standoff, the second term of TACE collapses toward zero, because residual penetration is then at or above its maximum and exceeds any realistic thin-roof threshold; protection rests almost entirely on the disruption term. When the gap is kept short, below the optimal standoff, so the jet has not yet reached full stretch, or pushed well beyond it onto the declining branch, the second term contributes meaningfully. Across the design space, modelled TACE ranges from about 0.28 for a coarse 60 mm grid placed near the optimal standoff to about 0.83 for a fine 30 mm grid at a sub-optimal short gap. Table 3 reports representative values.

Grid pitch (mm)	Roof gap (mm)	Standoff regime	P_d	TACE
30	250	sub-optimal (short)	0.75	0.83
30	350	near optimum	0.75	0.77
40	250	sub-optimal (short)	0.60	0.71
40	350	near optimum	0.60	0.62
50	450	declining branch	0.45	0.58
60	350	near optimum	0.34	0.28

Table 3. Illustrative TACE values across the cage design space for a representative ≈ 85 mm warhead against a thin tank roof. All values are model-derived and require experimental validation before any fielding decision.

The pattern is consistent and, for the design question, decisive: effectiveness is dominated by the disruption term, modulated by where the roof gap sits on the standoff curve, and worst when a coarse grid is combined with a gap near the warhead's optimal standoff. The analytical sections that follow explain why.

THE STANDOFF PARADOX IN TOP-ATTACK ENGAGEMENTS

The most consequential finding is also the least intuitive, and it deserves to be stated without hedging before it is qualified: adding standoff can increase penetration. The shaped-charge jet is not a rigid rod that loses energy with distance. It is a stretching column whose penetrating capacity grows as it elongates, up to the point where velocity gradients tear it into discrete particles that begin to tumble and drift off the firing axis (Elek et al., 2024; Jia et al., 2021). Before that point, every additional centimetre of standoff buys penetration; after it, every additional centimetre costs penetration. The curve therefore has an interior maximum, and the validated PG-7VM data locate that maximum near 360 mm of standoff, where penetration reaches about 317 mm against roughly 278 mm at contact (Żochowski et al., 2021). A naive roof cage that holds the warhead at, say, 300–360 mm from the roof does not move the engagement onto the safe, declining branch of this curve. It moves it toward the peak.

This is the precise sense in which transplanting slat armour from flank to roof imports hardware without physics. On the flank, the RPG round is defeated chiefly because the bars assault the warhead before a jet forms; the standoff change is incidental and, in any case, the round's optimal standoff was tuned for its own internal geometry, not for the bar spacing (Coghe, 2022). On the roof, against an FPV-delivered charge that may have been removed from its original launcher and re-fused, the cage's added gap interacts directly with the munition's standoff–penetration curve. If the operator's charge happens to be one whose optimum lies near the cage gap, premature initiation at the cage plane, exactly the outcome a cage is supposed to induce, can deliver a better-formed, longer jet to the roof than a contact detonation would. The disruption mechanism must therefore do the protective work; the standoff change is at best neutral and at worst counterproductive.

The parametric studies make the magnitude tractable. Dehestani et al. (2019) show that decreasing liner thickness and tuning standoff jointly govern penetration efficiency, and that the penetration-versus-standoff relation is smooth and single-peaked across the practical range. Elek et al. (2024) reproduce the same qualitative behaviour from a virtual-origin model complemented by Eulerian simulation, linking the optimum to the jet's effective standoff and velocity gradient. The convergence of an experimental campaign (Żochowski et al., 2021), a parametric numerical study (Dehestani et al., 2019), and an analytical–numerical model (Elek et al., 2024) on the same non-monotone shape is what gives the standoff paradox its evidentiary weight. It is not an artefact of one dataset.

Two design implications follow, and they are sharper than the usual call for 'more standoff'. The first is that a roof cage should keep its gap demonstrably short, below the warhead's optimal standoff, so that any jet which does form is caught while still under-stretched, rather than at or beyond its peak. A short gap also happens to be the easier retrofit, because it intrudes less on hatch operation and vehicle signature. The second is that if a long gap is chosen at all, it must be long enough to push the engagement well onto the declining branch, which for the modelled warhead means substantially more than 360 mm and quickly becomes impractical on a turret roof. Between these two regimes lies a trough of poor performance centred on the optimum, and it is exactly where an unthinking 'just add a cage' retrofit tends to land. The TACE index makes that trough visible: in Table 3, the near-optimum rows carry the lowest composite values at every pitch.

I want to be careful not to overclaim the generality of the 360 mm figure. It is specific to the modelled warhead and its liner; a smaller custom charge improvised onto an FPV will have a proportionally smaller optimal standoff, perhaps well under 200 mm, while a larger tandem-style charge will sit higher. The optimal standoff scales with the warhead's cone diameter and should be rescaled for any other munition (Żochowski et al., 2021). What generalises is the existence and approximate location, a few charge diameters, of the optimum, not its absolute value.

The standoff paradox also reframes a common observation from the field, namely that cages 'sometimes work and sometimes do not.' On the present analysis that inconsistency is not mysterious. It is the expected behaviour of a device whose protective value depends on whether the operator's particular charge, fused and aimed in a particular way, is disrupted by the grid, and, failing disruption, on whether the cage gap happens to fall on the rising, peak, or declining branch of that charge's standoff curve. The same cage can sit in different regimes against different munitions, which is why single anecdotes neither condemn nor vindicate the concept. Only the joint geometry does.

FUZE INTERACTION AND PREMATURE INITIATION ON A ROOF-MOUNTED GRID

If disruption carries the protective burden, then the engineering question narrows to how reliably a roof grid can insult an FPV-borne warhead before it forms a jet. The slat-armour literature gives a clear mechanism and a sobering caveat. The mechanism is geometric: protection is the probability that the warhead's fuze, or the body immediately behind it, strikes a grid element rigid enough to crush, short, or shear it, rather than threading the aperture between bars (Coghe, 2022). The caveat is that this probability is a strong function of impact obliquity, and that net and lattice systems lose efficiency rapidly once the impact angle climbs away from the perpendicular flank case, Coghe (2022) reports a marked drop above roughly 15 degrees of vertical angle. A top-attack FPV arrives near the worst angle for the mechanism the cage relies upon.

Three features of the FPV threat make the disruption problem distinctive rather than merely a re-run of the RPG case. The terminal velocity is low: an FPV quadcopter typically closes at some tens of metres per second, far below the 100–290 m/s of a boosted RPG round, which changes the energy available to crush the warhead against a bar and lengthens the dwell time during which a piezoelectric or mechanical fuze may or may not initiate. The warhead is often improvised: charges are removed from RPG, RKG, or other munitions and re-fused with contact or barometric triggers whose stiffness and standoff differ from the original (Kweera, 2023; Mittal & Goetz, 2025). And the aim point is chosen: an operator watching a video feed can steer the dive toward a hatch ring, a cupola, or the gap between cage segments, defeating a lattice that protects only the area it covers (Teixeira da Silva, 2023). Each feature degrades the clean fuze-strike geometry that the slat model assumes.

Premature initiation, when it does occur, is the most favourable outcome, but only on the conditions the previous section established. If the grid initiates the warhead at a short standoff, below the optimum, the jet reaches the roof under-stretched and with reduced penetration, and the protection term of TACE rises accordingly. If the grid merely strips a fuze or bends a fin without initiating, the warhead may pass through and detonate on the roof at near-contact standoff, which for the modelled charge corresponds to the 278 mm penetration of the rising branch, still ample against a thin roof, but below the 317 mm peak (Żochowski et al., 2021). The least favourable disruption-free outcome is initiation at the optimum. The design intent, then, is not simply 'initiate early' but 'initiate early and short', which argues for a grid dense enough to force interaction and a gap shallow enough to keep any resulting jet under-developed.

The reactive- and active-protection comparison clarifies how modest the cage's mechanism is by contrast. Explosive reactive armour does not rely on a probabilistic fuze strike; it drives a plate across the jet to erode and deflect it deterministically, and modern optimisation can tune cassettes to defeat both long-rod and shaped-charge threats at minimum mass (Mayseless et al., 2019; Moldtmann et al., 2024). Even against explosively formed projectiles, where the interaction window is short, reactive layers impose measurable penetration loss (Nie et al., 2019). Active systems remove the probabilistic dependence on the threat striking a particular element by engaging the munition in flight, with ensemble-learning controllers now timing the countermeasure against a moving platform (Li et al., 2023) and combat-effectiveness models quantifying their kill probability (Park et al., 2024). The cage, by comparison, offers a single probabilistic mechanism at the least favourable obliquity. That it works at all is a function of how many FPV warheads are crude and how many operators misjudge the aim point, contingencies, not guarantees.

There is a further subtlety that the disruption framing should not obscure. Even a non-initiating interaction that deforms the liner or tips the warhead off axis degrades the jet that subsequently forms, because jet coherence is exquisitely sensitive to initiation symmetry and liner integrity (Du et al., 2021; Pyka et al., 2020). A warhead that detonates while yawed delivers a jet whose virtual origin is displaced and whose penetration along the roof normal is reduced. This 'soft' disruption is hard to credit precisely in a single probability, and in the model it is folded conservatively into P_d, but it is one reason field results sometimes exceed what a pure fuze-strike calculation would predict. I treat it as a source of upside that the index deliberately understates rather than as a mechanism to be relied upon.

MODULAR GEOMETRY, MASS BUDGETS, AND OPERATIONAL TRADE-OFFS IN TRENCH WARFARE

A protective concept that survives the physics must still survive the vehicle and the war. Here the trench-warfare setting bites hardest, because it removes the mitigations a moving vehicle enjoys and amplifies the costs an add-on imposes. A vehicle that loiters cannot use speed to spoil the terminal dive, presents a stationary aim point for minutes or hours, and is engaged repeatedly rather than once (Mittal & Goetz, 2025). Any roof structure must therefore be judged not by its single-shot index alone but by its mass, its effect on mobility and hatch operation, its repairability under field conditions, and its behaviour under multiple strikes. The case for modularity follows directly from these constraints.

Mass is the first-order constraint, and survivability research insists on treating it as such. Add-on protection competes for the same weight budget as mobility, and crew-protection studies show that mass added high on the vehicle degrades stability and the very agility that helps against other threats (Pratomo et al., 2020; Kciuk et al., 2022). The slat-armour mass-efficiency results are instructive: polymer-fibre net systems achieve the best protection per kilogram, but their efficiency degrades steeply at the vertical impact angles a roof cage faces (Coghe, 2022). The designer is thus pushed toward the lightest grid that still delivers an adequate disruption probability, favouring fine-pitch, low-mass lattices over heavy bar arrays, precisely because Table 2 shows that pitch, not bar mass, drives P_d. A 30 mm fine grid in a light alloy or composite can outperform a heavier 60 mm steel bar array on both mass and effectiveness.

Modularity addresses the multi-strike and repair problem that a monolithic cage cannot. In trench conditions a roof cage will be hit, damaged, and hit again; a welded one-piece structure that is locally destroyed leaves a hole that the operator's camera will find. A panelised, boltable system lets a crew replace a damaged segment in minutes from within or beside the protected volume, an operational logic that mirrors the reload-under-armour philosophy now emphasised for contested environments. The same modularity allows the geometry to be tuned per threat axis: denser panels over the most-attacked turret roof and hatch rings, sparser or absent panels where intrusion on optics and weapons would cost more than it protects. This is the practical content of the second hypothesis, an optimal window exists, and modular construction is what lets a unit actually occupy it rather than admire it.

The geometry interacts with the rest of the vehicle in ways a roof-only view misses. A cage tall enough to push the gap onto the declining branch of the standoff curve, well beyond 360 mm for the modelled charge, would foul the main gun's elevation arc, obstruct the commander's sight, and enlarge the radar and visual signature that cues the attack in the first place (Castrillo et al., 2022). A short cage avoids these penalties and, as the standoff analysis showed, is the regime that keeps any jet under-stretched. The convergence is fortunate: the geometry that is best for protection against the modelled charge — short gap, fine pitch — is also the one that least degrades mobility, signature, and weapon employment. Where this convergence fails, it fails because the operator's charge has an unusually short optimum, in which case no practical gap can outrun it and disruption must do all the work.

Even granting an optimal modular cage, the trench setting imposes a residual that passive structure cannot remove: the chosen aim point. An operator who steers into the gap between panels, into an open hatch, or into the engine deck behind the cage circumvents the protected zone entirely (Teixeira da Silva, 2023). Coverage is therefore a design variable as important as pitch and gap, and complete coverage is impossible without immobilising the vehicle. This is the structural reason the third hypothesis must hold: the cage, however well designed, leaves uncovered axes and a probabilistic disruption mechanism, and the residual it leaves can only be closed by other layers. The cage buys time and probability, not certainty, and in a war of repeated engagement probability compounds against the defender unless other tiers shorten the exposure.

CAGE ARMOUR WITHIN THE COUNTER-UAS SURVIVABILITY HIERARCHY

The defensible role for the cope cage emerges once it is placed where it belongs, at the outer edge of a graduated survivability stack rather than at its centre. The counter-UAS literature describes a layered option space whose tiers engage the threat at decreasing range: detection and soft-kill denial of the radio-frequency or satellite-navigation link, hard-kill interception by directed energy or cooperative defensive drones, and only then the terminal, on-vehicle measures (Castrillo et al., 2022; Lee et al., 2022; Silva et al., 2023). Each tier has a characteristic failure mode. Soft-kill assumes a link to break, yet terminally autonomous FPVs that lock on before jamming can still strike (Mittal & Goetz, 2025; Kweera, 2023). Hard-kill assumes detection and reaction time that a low, fast, late-unmasking FPV may deny. The terminal tier exists precisely to catch what the outer tiers miss.

Within the terminal tier, the cage is the cheapest and most degraded option, and that is an argument for it, not against it, provided it is not asked to work alone. Active protection systems engage the munition before impact and remove the cage's dependence on a probabilistic fuze strike, but they are costly, finite in magazine, and historically optimised against horizontal threats, with top-attack coverage only now maturing (Li et al., 2023; Park et al., 2024). Reactive armour deterministically disrupts a jet but adds mass and is itself a hazard to nearby infantry, a serious constraint in the combined-arms reality of trench fighting (Mayseless et al., 2019; Moldtmann et al., 2024; Nie et al., 2019). The cage occupies the niche beneath all of these: near-zero unit cost, field-repairable, and effective against the substantial fraction of FPV warheads that are crude enough to be disrupted by a grid. It converts some lethal hits into mobility kills or misses, and in an economy where the attacker spends hundreds of dollars and the defender millions, even a partial conversion alters the exchange ratio that drives the campaign (Mittal & Goetz, 2025).

The systems view also disciplines expectations about what the cage can and cannot do. It cannot defeat a precisely aimed strike into an uncovered aperture, a tandem charge designed to defeat standoff, or an operator who simply selects a different impact point; those are the responsibilities of the inner tiers and of tactics, dispersion, overhead concealment, and the denial of loiter time (Castrillo et al., 2022). What it can do is raise the probability that a generic, mass-produced FPV with an improvised warhead fails to achieve a clean perforation, and do so at a cost low enough to apply across an entire fleet rather than a privileged few vehicles. That fleet-wide applicability is the cage's real comparative advantage over active and reactive systems, which budgets ration to a minority of platforms.

Read together, the four analytical strands answer the research question with more precision than the binary 'do cages work' allows. A roof cage measurably lowers perforation probability when its pitch is fine enough to drive a high disruption probability and its gap is short enough to keep any surviving jet under-stretched; it fails, or backfires, when a coarse grid is paired with a gap near the warhead's optimal standoff; and even at its optimum it is a probabilistic outer layer whose value is realised only inside a graduated counter-UAS architecture. The TACE index is the instrument that turns those qualitative statements into a design space a unit can navigate.

CONCLUSION

The cope cage deserves neither the ridicule of its detractors nor the faith of its improvisers, and the reason is that its value is conditional on geometry that most field retrofits do not deliberately choose. This article set out to determine when a roof-mounted modular cage measurably protects a main battle tank against a top-attack, FPV-delivered shaped charge, and the analysis supports a graded rather than a categorical answer.

The first hypothesis, that disruption, not standoff, is the dominant protective mechanism, and that a standoff-only design can increase penetration, is confirmed. The standoff–penetration curve has an interior maximum near 360 mm for the modelled warhead, where penetration reaches about 317 mm against 278 mm at contact (Żochowski et al., 2021), and three independent lines of evidence converge on that non-monotone shape (Dehestani et al., 2019; Elek et al., 2024; Jia et al., 2021). A cage that holds an undisrupted warhead near that optimum can therefore deliver a better-formed jet to the roof than contact detonation would, and the protective work must be done by warhead disruption, whose probability the slat-armour methodology lets us estimate (Coghe, 2022).

The second hypothesis, that a bounded, computable geometric window maximises effectiveness, is confirmed and given quantitative content. Modelled TACE rises above 0.80 for a fine 30 mm grid at a short, sub-optimal gap and falls below 0.30 for a coarse 60 mm grid placed near the optimal standoff, with a trough of poor performance centred on the optimum at every pitch. The window that a designer should occupy is fine pitch and short gap, and that window happens to coincide with the lightest, least intrusive, most repairable retrofit, an alignment of protective and operational interests that modular construction is well suited to exploit.

The third hypothesis, that even an optimal cage is insufficient alone, is confirmed without qualification. The cage leaves uncovered axes, depends on a probabilistic mechanism at the least favourable obliquity, and faces an operator who chooses the aim point; in a war of repeated engagement those residuals compound. Its defensible role is as the outermost, cheapest, fleet-wide tier of a layered counter-UAS architecture that also fields soft-kill, hard-kill, and reactive elements (Castrillo et al., 2022; Li et al., 2023; Moldtmann et al., 2024; Mittal & Goetz, 2025).

The principal original contribution of this article is the Top-Attack Cage Effectiveness index, which joins the grid-disruption probability of slat-armour research to the standoff-dependent residual penetration of terminal-ballistics research to define, for the first time to my knowledge, an optimal roof-cage geometry for the top-attack FPV case. The index is deliberately simple, and its simplicity is the point: it gives a unit a way to reason about pitch and gap together rather than chasing standoff alone, and it makes the standoff paradox, the regime in which a cage helps the attacker, explicit and avoidable.

The limitations are real and I have tried to keep them in view rather than in a footnote. The disruption probability is borrowed from an RPG-validated model and is indicative for re-fused FPV warheads; the standoff curve is calibrated to a single well-documented charge and will shift for others, though its non-monotonicity is general; and the index treats one engagement, not the repeated attacks that define trench warfare. Each is a boundary on the claims, not a refutation of them. The natural next step is experimental: a proof-range campaign that fires representative improvised warheads against instrumented modular grids of varied pitch and gap would convert the model-derived TACE values into measured ones and, I suspect, sharpen rather than overturn the design rules proposed here. Until that campaign is run, the numbers in this article should guide thinking, not procurement.

A closing reflection from the vantage of defence engineering rather than tactics. The cope cage is an unusually honest artefact: it is what a crew builds when the threat outpaces the institution, and its very crudeness has forced a re-examination of standoff physics that the armoured community had treated as settled. The contribution here is to give that improvisation a model worthy of it, so that the next grid welded onto a turret is shaped by the standoff curve and the disruption geometry rather than by hope.

BIBLIOGRAPHY

Becker, M., Klavzar, A., Wolf, T., & Renck, M. (2022). Data-driven prediction of plate velocities and plate deformation of explosive reactive armor. *Defence Technology, 18*(12), 2141–2149. https://doi.org/10.1016/j.dt.2022.07.001

Castrillo, V. U., Manco, A., Pascarella, D., & Gigante, G. (2022). A review of counter-UAS technologies for cooperative defensive teams of drones. *Drones, 6*(3), 65. https://doi.org/10.3390/drones6030065

Chen, J.-Y., Feng, D.-L., Sun, Q.-Y., Peng, C., Zhu, Y.-H., & Yu, S.-Y. (2023). Numerical modeling of shaped charge jet penetration into ceramic–metal double-layered medium using smoothed particle hydrodynamics. *International Journal of Impact Engineering, 175*, 104526. https://doi.org/10.1016/j.ijimpeng.2023.104526

Coghe, F. (2022). Efficiency of different cage armour systems. *Applied Sciences, 12*(10), 5064. https://doi.org/10.3390/app12105064

Dehestani, P., Fathi, A. R., & Mohammadi Daniali, H. R. (2019). Numerical study of the stand-off distance and liner thickness effect on the penetration depth efficiency of shaped charge process. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science, 233*(3), 977–986. https://doi.org/10.1177/0954406218766183

Du, Y., He, G., Liu, Y., Guo, Z., & Qiao, Z. (2021). Study on penetration performance of rear shaped charge warhead. *Materials, 14*(21), 6526. https://doi.org/10.3390/ma14216526

Du, Y., He, G., Li, W., & Wang, K. (2022). Experimental and numerical study on the penetration performance of a shaped charge. *Materials, 15*(11), 3899. https://doi.org/10.3390/ma15113899

Elek, P., Marković, M., Jevtić, D., & Đurović, R. (2024). Modeling of shaped charge jet penetration depth: Analytical and numerical approach. *Scientific Technical Review, 74*(2), 17–22. https://doi.org/10.5937/oteh2402017e

Elshenawy, T., Li, Q. M., & Elbeih, A. (2022). Experimental and numerical investigation of zirconium jet performance with different liner shapes design. *Defence Technology, 18*(1), 12–26. https://doi.org/10.1016/j.dt.2020.11.019

Elshenawy, T., Abdo, G. M., & Elbeih, A. (2024). Effect of explosives charge types on the jet characteristics, penetration performance and fragmentation patterns of shaped charges. *Scientific Reports, 14*, 26282. https://doi.org/10.1038/s41598-024-75727-0

Hao, Z., Wang, Z., Wang, Y., Duan, C., & Ji, Q. (2023). The effect of three-layer liner on the jet formation and penetration capability of shaped charge jet. *Scientific Reports, 13*, 13851. https://doi.org/10.1038/s41598-023-38680-y

Jia, X., Xu, M., Huang, Z., Xiao, Q., Chen, S., & Ma, B. (2021). Theoretical analysis and experimental study of the performance of shaped charge jet penetration into thick-walled moving target by rocket sled testing method. *International Journal of Impact Engineering, 155*, 103894. https://doi.org/10.1016/j.ijimpeng.2021.103894

Kciuk, S., Krzystała, E., Mężyk, A., & Szmidt, P. (2022). The application of microelectromechanical systems (MEMS) accelerometers to the assessment of blast threat to armored vehicle crew. *Sensors, 22*(1), 316. https://doi.org/10.3390/s22010316

Kweera, R. (2023). Drones in modern warfare: Utilization in India–Pakistan cross-border terrorism and security implications. *Strategic Analysis, 47*(4), 376–388. https://doi.org/10.1080/09700161.2023.2288989

Lee, C. H., Thiessen, C., Van Bossuyt, D. L., & Hale, B. (2022). A systems analysis of energy usage and effectiveness of a counter-unmanned aerial system using a cyber-attack approach. *Drones, 6*(8), 198. https://doi.org/10.3390/drones6080198

Li, C., Wang, G., Song, H., Huang, X., & Zhou, Q. (2023). Angular disturbance prediction for countermeasure launcher in active protection system of moving armored vehicle based on an ensemble learning method. *Defence Technology, 21*, 207–218. https://doi.org/10.1016/j.dt.2022.10.007

Ma, T., Liu, J., & Wang, Q. (2022). Structural design and parameter optimization of linear explosive forming projectile. *Propellants, Explosives, Pyrotechnics, 47*(4), e202100348. https://doi.org/10.1002/prep.202100348

Mayseless, M., Friling, S., & Misiuk, L. (2019). Re-visiting the mass-flux model for explosive reactive armor and the effect of plate thickness. *Defence Technology, 15*(5), 779–785. https://doi.org/10.1016/j.dt.2019.08.015

Mittal, V., & Goetz, J. (2025). A quantitative analysis of the effects of drone and counter-drone systems on the Russia–Ukraine battlefield. *Defense & Security Analysis, 41*(3), 490–503. https://doi.org/10.1080/14751798.2025.2479973

Moldtmann, P., Berk, J., Ryan, S., & Klavzar, A. (2024). Adaptive optimisation of explosive reactive armour for protection against kinetic energy and shaped charge threats. *Defence Technology, 40*, 1–12. https://doi.org/10.1016/j.dt.2024.05.007

Nie, Y., Jiang, J., Wang, S., Liu, H., & Men, J. (2019). Penetration capability of EFPs against explosive reactive armor. *Propellants, Explosives, Pyrotechnics, 44*(5), 597–608. https://doi.org/10.1002/prep.201800187

Park, S., Park, J., Suh, S., & Kim, J. (2024). Analysis of combat effectiveness of active protection system of armored vehicles. *Journal of the Korea Academia-Industrial Cooperation Society, 25*(7), 423–432. https://doi.org/10.5762/KAIS.2024.25.7.423

Pratomo, A. N., Santosa, S. P., Gunawan, L., & Putra, I. S. (2020). Countermeasures design and analysis for occupant survivability of an armored vehicle subjected to blast load. *Journal of Mechanical Science and Technology, 34*(5), 1893–1899. https://doi.org/10.1007/s12206-020-0411-1

Pyka, D., Kurzawa, A., Bocian, M., Bajkowski, M., Magier, M., Sliwinski, J., & Jamroziak, K. (2020). Numerical and experimental studies of the ŁK type shaped charge. *Applied Sciences, 10*(19), 6742. https://doi.org/10.3390/app10196742

Shi, J., Huang, Z., Zu, X., & Xiao, Q. (2022). Experimental and numerical investigation of jet performance based on Johnson–Cook model of liner material. *International Journal of Impact Engineering, 170*, 104343. https://doi.org/10.1016/j.ijimpeng.2022.104343

Silva, D. L. D., Machado, R., Coutinho, O. L., & Antreich, F. (2023). A soft-kill reinforcement learning counter unmanned aerial system (C-UAS) with accelerated training. *IEEE Access, 11*, 31496–31507. https://doi.org/10.1109/ACCESS.2023.3253481

Teixeira da Silva, J. A. (2023). Can a tank 'umbrella' protect the turret of a main battle tank? *Scientific Journal of the Military University of Land Forces, 210*(4), 361–366. https://doi.org/10.5604/01.3001.0054.1632

Yang, D., & Lin, J. (2021). Numerical investigation on the formation and penetration behavior of explosively formed projectile (EFP) with variable thickness liner. *Symmetry, 13*(8), 1342. https://doi.org/10.3390/sym13081342

Zhang, H., Zheng, Y., Yu, Q., Ge, C., Su, C., & Wang, H. (2022). Penetration and internal blast behavior of reactive liner enhanced shaped charge against concrete space. *Defence Technology, 18*(6), 952–962. https://doi.org/10.1016/j.dt.2021.04.011

Żochowski, P., & Warchoł, R. (2023). Experimental and numerical study on the influence of shaped charge liner cavity filling on jet penetration characteristics in steel targets. *Defence Technology, 23*, 60–74. https://doi.org/10.1016/j.dt.2022.09.007

Żochowski, P., Warchoł, R., Miszczak, M., Nita, M., Pankowski, Z., & Bajkowski, M. (2021). Experimental and numerical study on the PG-7VM warhead performance against high-hardness armor steel. *Materials, 14*(11), 3020. https://doi.org/10.3390/ma14113020

MODULARNI ANTIDRONSKI KAVEZNI OKLOP NA GLAVNIM BORBENIM TENKOVIMA: INŽENJERSKA PROCJENA EFIKASNOSTI PROTIV KUMULATIVNIH BOJEVIH GLAVA KOJE NOSE FPV DRONOVI U USLOVIMA ROVOVSKOG RATOVANJA

Emre Aydin

Agencija za odbrambenu industriju
Ankara, Republika Turska
E-pošta: emre.aydin@ssb.gov.tr

Originalni naučni članak

Primljeno: 17.03.2026. Prihvaćeno: 21.05.2026.

DOI: https://doi.org/10.65932/military-studies-2026-1-2 **UDK**: 623.438.3:623.562

Sažetak: Dronovi s prvoličnim pogledom (FPV) koji nose prenamijenjene kumulativne bojeve glave preokrenuli su logiku troška-iscrpljivanja u mehanizovanom ratovanju, a posade su odgovorile improvizovanim krovnim rešetkama, 'kavezima', čija inženjerska vrijednost ostaje sporna. Ovaj rad postavlja usko pitanje: pod kojim geometrijskim uslovima krovni modularni kavez mjerljivo smanjuje vjerovatnoću proboja kumulativne bojeve glave koju FPV dron isporučuje napadom odozgo, a kada to ne uspijeva? Istraživanje povezuje model vjerovatnoće interakcije s rešetkom, razvijen za rešetkasti oklop protiv RPG-7 prijetnje, s empirijski utemeljenim odzivom prodiranja u funkciji odstojanja za reprezentativnu bojevu glavu i s utvrđenom mehanikom formiranja i prodiranja mlaza. Originalni doprinos rada jeste indeks efikasnosti krovnog kaveza pri napadu odozgo (TACE), bezdimenzionalna kompozitna mjera koja množi vjerovatnoću ometanja bojeve glave rešetkom s preostalom marginom prodiranja mlaza koji preživi inicijaciju na odstojanju kaveza, vrednovanom u odnosu na prag zaštite krova. Indeks čini eksplicitnim kontraintuitivan rezultat: pošto kumulativni mlaz dostiže maksimum prodiranja na konačnom optimalnom odstojanju, oko 317 mm naspram 278 mm pri kontaktu za modeliranu bojevu glavu, kavez koji samo dodaje odstojanje, bez ometanja bojeve glave, može povećati prodiranje umjesto da ga smanji. Ilustrativni proračun za korake rešetke od 30 do 60 mm i krovne razmake od 250 do 450 mm pokazuje vrijednosti TACE od ispod 0,30 za loše postavljene grube rešetke do iznad 0,80 za fine konfiguracije u kojima dominira ometanje. Zaključuje se da su kavezi branjivi kao najvanjskiji sloj postupne protivdronske arhitekture preživljavanja, a ne kao samostalna zaštita, i da njihov dizajn mora voditi geometrija ometanja bojeve glave, a ne samo odstojanje. Sve numeričke vrijednosti TACE su izvedene iz modela i zahtijevaju eksperimentalnu validaciju prije primjene.

Ključne riječi: *kavezni oklop, FPV dronovi, kumulativna bojeva glava, odstojanje, zaštita od napada odozgo, glavni borbeni tenk, preživljavanje.*