
Quantifying Layered Naval Defense Against Hypersonic Glide Vehicles

A Monte Carlo Simulation Study of SM-6 and PAC-3 MSE
Terminal Defense Against DF-17 HGV Salvos

1,500 Monte Carlo Runs | 3 Loadout Variants | 4 Evasion Profiles

Empyrean Defense Research Brief - April 2026

UNCLASSIFIED // OPEN SOURCE INTELLIGENCE ONLY



DATA SOURCES & METHODOLOGY

This publication uses publicly available data from open-source providers including CSIS Missile Threat Project, The War Zone, Navy FY2027 budget justification documents, Lockheed Martin published specifications, and peer-reviewed academic publications. No classified or restricted data was used at any point in this research.

Empyrean Defense has no commercial relationship with any defense contractor, government agency, or weapons manufacturer referenced in this document.

Physics models used in this simulation are drawn from published, unclassified sources including Kingery-Bulmash blast overpressure polynomials (BRL-TR-02555), Zarchan proportional navigation guidance, Tsiolkovsky rocket equation, ISA atmospheric model, and UFC 3-340-02 structural response criteria. All models carry inherent simplification relative to high-fidelity computational methods; results should be treated as analytical approximations, not operational predictions.

All entity parameters (interceptor specifications, threat characteristics, ship properties) are derived from open-source intelligence and may not represent actual operational performance. Results represent upper-bound performance estimates under defender-favorable assumptions.



About Empyrean Defense Research

Empyrean Defense Research produces unclassified, physics-backed analysis of real wargaming scenarios across every domain in Joint All-Domain Operations (JADO). Our intent is to be useful - from Congressional staffers evaluating program funding to the depot technician maintaining the effector - not to propagandize results. As Americans, we may not always like the findings. Physics doesn't care.

The Empyrean Defense Wargaming & Simulation Cyber Range began as an organic training module inside our platform - the Decision Dominance Engine (DDE) - and grew into a high-fidelity simulation engine that models sea states, RF propagation, CEC uplinks, beamforming, blast pressure, thermoclines, atmospheric extinction, and more across dozens of public-domain, widely-accepted physics formulas. We cannot account for every variable in a live engagement, and we err on the side of optimism from detection through battle damage assessment. Where we approximate, we say so.

Every entity in our scenarios - what we call a "playing card" - is built from open-source, public-domain data; aggregated from U.S. and allied defense publications, academic research, OSINT blogs, manufacturer specifications, and social media. We do not claim these numbers will survive the rigors of a real fight. We do not claim to know the exact RCS nor fuze proximity parameter, or anything else other than the manufacturers and maintainers would know.

We do claim the math is honest, the methodology is transparent, and every simulation run is reproducible from the scenario file and seed alone. In the interest of protecting our core IP, we cannot make that available, as it runs against a live engine.

Foreword

I want to be honest about what this research is and what it is not.

We ran 1,500 Monte Carlo simulations of an Arleigh Burke Flight IIA destroyer defending itself against PRC DF-17 hypersonic glide vehicle salvos. We tested three defensive loadout variants - including the PAC-3 MSE/Aegis integration that the Navy plans to field by FY2027 - across four evasion profiles and two salvo sizes. We gave the defender every advantage we could think of: perfect radar, perfect guidance, perfect crew, no jamming, no multi-axis attack, calm seas, clear weather. We stacked the deck for the good guys.

The ship died in every run. Every variant. Every evasion level. Every time.



PAC-3 MSE/Aegis that will be fitted for Mk 41 VLS aboard an Arleigh achieved zero kills across 450 terminal resolution events, no matter if, or how fast, the HGV was jinking on terminal approach. The hit-to-kill guidance designed for Mach 3-4 theater ballistic missiles cannot close the miss distance gap at Mach 10 closing velocity. The physics do not care about the integration timeline or the budget justification.

I am not a naval warfare officer, the only vessel I've ever boarded was a cruise ship or many charter fishing boats in the Long Island Sound. I am a defense software engineer who builds sensor fusion and decision support tools. What follows is a physics-grounded simulation study under optimistic assumptions, not an operational assessment, I'm not qualified to do so. The findings should be treated as a floor, not a ceiling: every factor we did not model (electronic attack, crew stress, maintenance readiness, multi-axis coordinated salvo) makes the real number worse.

If you are a program office evaluator, a surface warfare officer, or a Congressional staffer reading this: **the engagement timeline is the problem**. Not the interceptor nor the seeker nor the warhead. The timeline. Even if our drag model is off by 20%, or our speed regimes are off by 50%, the results would likely be the same. Trust me, I tried to cook parametrics, I agonized about publishing this - no less for our FIRST Empyrean Defense Research project - I want American warfighters to win and come back to the things they love.

Oh, and that timeline? It is purely geometry, not engineering. You cannot so easily defeat geometry with a better missile.

What you can do about it is addressed in the final section. There are answers. They are not cheap and they are not fast, but they exist.

- Empyrean 7

Abstract

In the interest of minimizing as many variables as we can (Electronic Attack, decoys, layered threat types per salvo, multiple ships, CEC bandwidth, and more) we placed a single Arleigh Burke Flight IIA destroyer off the Second Island Chain, faced with DF-17 Hypersonic Glide Vehicles (HGVs) on terminal approach. This is under the assumption that from a reported DF-17 Transport, Erector, Launcher (TEL) magazine of 12, not all of the DF-17s will survive to their separation state where they release their HGV on final.

Across 1,500 Monte Carlo engagements (3 defensive loadout variants x 4 evasion profiles x 2 salvo sizes x 50 iterations per combination), a single Arleigh Burke Flight IIA destroyer failed to survive a DF-17 hypersonic glide vehicle salvo under any tested configuration. The ship was destroyed at t=139s in 100% of runs across all



variants. This is not an aberration, we placed the origin of each HGV just outside of engagement distance to give the sensors - which are also faithfully modeled with all the physics knowable to us - time to detect and engage.

PAC-3 Missile Segment Enhanced (MSE)/Aegis interceptors, planned for Mk 41 VLS integration by FY2027, achieved zero kills across all evasion levels - including non-maneuvering ballistic reentry (0G). The hit-to-kill mechanism fails at hypersonic closing velocities regardless of target maneuver profile. Open source and journalists point to the latest versions having a secondary fragmentation mechanism that disperses titanium fragmentation, that did not make a difference no matter if we decided to eschew Inverse Square Law and make them more lethal.

SM-6 Block I blast-fragmentation interceptors achieved 40-74% posthumous intercept rates (kills occurring after ship destruction) against staggered arriving HGVs, with performance paradoxically improving against maneuvering targets due to their extended time-in-zone executing the jinks.

Three findings challenge prevailing assumptions about hypersonic defense:

- 1. The engagement timeline, not interceptor performance, is the binding constraint.** The ship is destroyed before any interceptor reaches terminal phase. For the most realistic G-maneuvers that a DF-17 HGV will pull (10G) the average was 2.7 kills across the runs, but those are HGVs still approaching the ship which is already destroyed.
- 2. Hit-to-kill fails at hypersonic closing velocities even against non-maneuvering targets.** The failure mode is guidance bandwidth, not target evasion.
- 3. Terminal evasion maneuver helps the defender, not the attacker.** Maneuvering extends time-in-zone for blast-fragmentation interceptors. The optimal attacker strategy against SM-6 is a straight-line ballistic approach at maximum speed.

These results represent upper-bound performance estimates. Every modeling assumption favors the defender. If the physics alone constrains performance to these levels, the operational number is lower.

Data Provenance & Limitations

All entity parameters are derived from open-source intelligence: CSIS Missile Threat Project, The War Zone, Navy FY2027 budget justification documents, Lockheed Martin published specifications, and peer-reviewed academic publications on warhead lethality and blast mechanics. Source URLs are catalogued per entity card in the Empyrean Defense Platform repository. No classified or restricted data was used at any point in this



research.

Physics Models

The simulation resolves engagements through a deterministic physics pipeline, this is only representative of what was useful for this simulation.

Guidance. Three-dimensional proportional navigation using the Zarchan formulation. Navigation constant $N=3$ midcourse, $N=5$ terminal. Acceleration commands computed from line-of-sight rate, clamped to card-specified maximum acceleration authority. This is the same guidance law used in published SM-6 and PAC-3 literature - it is not a simplification.

Energy management. Analytical Tsiolkovsky boost phase (variable mass, constant thrust), Mach-dependent drag coefficients with ISA standard atmospheric density model, propellant mass fraction depletion. Interceptors launch at system-specific cold-launch velocities (SM-6: 30 m/s from VLS ejection), boost to burnout speed, then coast with aerodynamic drag decay. Loft-dive profiles for long-range engagements follow published SM-2/SM-6 engagement doctrine.

Evasion. Target maneuver modeled as random jink with 1.5-second hold intervals. Lateral acceleration applied perpendicular to velocity vector via cross-product decomposition. At 20G (196 m/s^2) sustained for 1.5 seconds, lateral displacement per jink is approximately 220 meters - more than sufficient to displace the threat outside the hit-to-kill engagement volume at hypersonic closing velocities.

Damage. Kingery-Bulmash blast overpressure polynomials with surface burst coefficients (Jeon, Kim & Han 2017). Confined blast amplification factor of 1.6x for pre-perforated targets per UFC 3-340-02. Inverse-square fragment density distribution per Chandela 1986. Dual-mode fuze evaluation: proximity detonation when miss distance falls within proximity radius, contact detonation on direct hit.

Structural damage. Per-subsystem damage accumulation across typed subsystems (structure, propulsion, guidance, payload) with independent integrity pools. Blast overpressure and fragmentation damage distributed across subsystems using function-specific attenuation multipliers. Target destroyed when any critical subsystem's structural integrity is fully depleted.

Optimistic Bias Framework

Every modeling assumption in this study favors the defender. This is by design - the objective is to establish the upper bound of terminal defense performance, not to predict operational outcomes. The actual number, we fear, is worse.

1. Perfect guidance. Zero sensor-to-shooter latency, continuous guidance updates, zero seeker track jitter. Real-world divert response lag (~200ms from seeker update to thruster command) and seeker noise at



hypersonic closing velocities would further degrade intercept probability.

2. Perfect detection. AN/SPY-6(V)1 acquires threats at maximum detection range with zero false tracks. No sea clutter, no atmospheric ducting, no multipath propagation.

3. No electronic attack. Threats approach in a clean electromagnetic environment - no standoff jamming (e.g., PLAN Type 055 shipboard EW suites), no chaff, no decoys. A realistic PLARF strike package includes dedicated EA aircraft or co-deployed jammers.

4. No VLS launch failures. Every cell fires on command. Fleet-wide VLS reliability is likely not 100% per cell.

5. Perfect crew. No decision latency under stress, no degraded watchstanding, no battle damage from prior salvos, no maintenance readiness gaps. We assumed all VLS cells are 100% combat-ready.

6. Single-axis attack. All threats arrive from one bearing. Multi-axis attack (the reported doctrinal norm for PLARF anti-ship operations) compresses the engagement timeline by splitting SPY-6 beam time across array faces.

7. No concurrent threat environment. No subsonic or supersonic anti-ship cruise missiles, no fighters, no torpedo threats, no electronic attack aircraft. The Burke's entire fire control capacity is dedicated to the single HGV salvo.

8. No engagement pacing constraints. The fire control loop assigns interceptors at the maximum rate the orchestrator can process - our physics engine can resolve complex engagement chains in nanoseconds in denser scenarios. Real-world salvo spacing, engagement assessment, and shoot-look-shoot or multi-missile shoot-shoot-look-shoot doctrine add seconds per engagement cycle.

Known Approximations

While we are not Naval officers, we do know that a single Arleigh-Burke would likely not be the norm. As mentioned earlier in this research piece, we wanted to eliminate all variables. A real engagement would have multiple ships and dozens of carrier-borne aircraft, space-based assets, and ground-based assets to help mitigate this threat. This scenario was modeled with the assumption that there will be leakers through a concentrated PLAN, PLAAF, and PLARF attack against an American Carrier Strike Group.

Furthermore, these exclusions bound the analysis scope, each represents a potential avenue for future research:

SM-3 Block IIA exo-atmospheric intercept during glide phase is not modeled. Scenarios begin with HGVs in terminal dive at 60 km altitude. Earlier-layer intercept opportunities, including SM-3-class engagements where geometry permits, would provide additional defensive layers before the terminal timeline collapses. HGVs complicate exo-atmospheric BMD because they fly depressed trajectories below classic midcourse intercept altitudes and can maneuver during glide phase.



CEC/NIFC-CA cooperative engagement is not modeled. Engage-on-remote via E-2D Advanced Hawkeye extends detection range and may provide 1 or 2 additional SM-6 engagement opportunities before the terminal timeline collapses.

Terrain masking is not modeled. This engagement scenario was modeled to take place at 25.221399, 126.631077, between Okinawa and Miyakojima in open water.

Human decision latency is not modeled. This engagement scenario was modeled in a purely autonomous, human-out-of-the-loop mode, targets were prosecuted nearly immediately where the first feasible engagement presented itself.

Scenario Design

Engagement Geometry

Open ocean, Western Pacific. A single Arleigh Burke Flight IIA DDG (9,700 tonnes full displacement) is positioned at 25.225 N, 126.649 E, heading 270 degrees at 15 knots. The DF-17 HGV salvo approaches from approximately 400 km east at 60,000m altitude, traveling at Mach 10 (3,400 m/s). Arrival is staggered: pairs of HGVs separated by 30 seconds, simulating sequential transporter-erector-launcher (TEL) firings from a single PLARF battery. These missiles were placed to the East (and not the West) to mitigate any terrain blocking effects from the terrain while also placing the ship in a likely engagement zone.

This geometry represents a head-on engagement with the SPY-6's primary array face - what we assumed is the most favorable detection geometry for the defender.

Environment

Conditions are uniformly favorable to the defender: Sea State 3 (Beaufort), daytime, clear visibility, 70% relative humidity, zero precipitation, zero wind. No environmental degradation applied to sensor detection range, weapon guidance accuracy, or seeker performance. This eliminates weather as a variable and isolates the physics of the engagement.

Defensive Loadout Variants

Total air-defense cell allocation is held approximately constant across variants to isolate the interceptor variable. Non-AD cells (Tomahawk land-attack, VL-ASROC anti-submarine) are not modeled.



Variant	Designation	Slot 1 (Outer)	Slot 2 (Mid)	Slot 3 (Terminal)	AD Cells
A	Baseline	24x SM-6 Block I	24x SM-2 Block IIIA	24x ESSM Block 2 (6 cells, quad-packed)	54
B	PAC-3 Sub	12x SM-6 Block I	24x PAC-3 MSE	24x ESSM Block 2 (6 cells)	42
C	PAC-3 Heavy	6x SM-6 Block I	36x PAC-3 MSE	24x ESSM Block 2 (6 cells)	48

Variant A represents the current Flight IIA loadout. SM-2 Block IIIA is a semi-active blast-fragmentation interceptor designed for air defense against aircraft and cruise missiles. It does not carry the theater ballistic missile (TBM) engagement capability - against hypersonic glide vehicles, 24 SM-2 cells are dead weight. Only the 24 SM-6 cells can engage this threat class.

Variant B replaces the SM-2 allocation with PAC-3 MSE on a 1-for-1 cell basis. This converts 24 dead cells into BMD-capable interceptors. This is the Navy's FY2027 integration plan as described in The War Zone and Navy budget justification documents.

Variant C maximizes PAC-3 MSE allocation at the expense of SM-6 outer-layer capacity and strike flexibility. This tests whether concentrating hit-to-kill interceptors improves outcomes against maneuvering HGVs.

Interceptor Specifications

As noted, these are derived from open-source, publicly available data sources and may not be truly representative of the true performance parameters.

Parameter	SM-6 Block I (RIM-174A)	PAC-3 MSE (MIM-104F)
Speed (cruise)	Mach 3.5 (1,190 m/s)	Mach 5 (1,700 m/s)
Max acceleration	300 m/s ² (30.6G)	400 m/s ² (40.8G)
Engagement range	370 km	60 km
Warhead	64 kg blast-frag (42 kg Comp-B fill)	73 kg hit-to-kill + lethality enhancer
Kill mechanism	Proximity fuze, ~25m lethal blast-frag radius	Direct body-to-body impact, 5m lethality enhancer radius
Terminal guidance	Active radar homing (AMRAAM-derived), F&F;	Ka-band active radar seeker, F&F;
Seeker acq. range	~40 km	~20 km



Parameter	SM-6 Block I (RIM-174A)	PAC-3 MSE (MIM-104F)
Terminal pulse motor	None	+200 m/s delta-V
Unit cost (per round)	\$4.5M (FY24)	\$4.05M (FY27)
Launch mass	~1,500 kg	~315 kg

Sources: CSIS Missile Threat Project, Lockheed Martin published specs, Navy FY2027 budget justification, The War Zone, DTIC ADA319957.

Threat Specifications

Card ID	Evasion	Terminal Accel	Maneuver	Rationale
df-17-hgv-0g	0G	None	Ballistic reentry	Pure speed isolation
df-17-hgv-10g	10G	98 m/s ²	Random jink, 1.5s	Conservative OSINT floor
df-17-hgv	20G	196 m/s ²	Random jink, 1.5s	Assessed operational
df-17-hgv-30g	30G	294 m/s ²	Random jink, 1.5s	Next-gen upper bound

All variants share identical warhead (500 kg unitary, 315 kg Comp-B equivalent), approach kinematics (Mach 10, 60 km ceiling), radar cross section (0.01 m² - low-observable glide body), and structural hardening (800 kPa blast overpressure threshold, consistent with reentry-grade thermal protection and airframe materials).

At Mach 10, the kinetic energy of a 1,400 kg HGV glide body at impact is approximately 8 gigajoules - roughly equivalent to 1.9 tonnes of TNT. The 315 kg Comp-B warhead adds another 1.3 GJ of explosive energy. Total energy delivery per HGV exceeds 9 GJ. For context, a Harpoon anti-ship missile delivers approximately 0.04 GJ. The DF-17 delivers over 200 times the energy of the weapon the Burke was originally designed to survive.

Salvo Sizes

4 and 8 HGVs tested as primary sweep axes. 12 and 16 HGVs were also tested but produced identical results to 8 due to salvo saturation - the ship is destroyed by the first HGV at t=139s, and subsequent arrivals impact after destruction (see Finding 4).

Sweep Matrix

4 evasion levels x 3 loadout variants x 2 salvo sizes = 24 scenario combinations x 50 Monte Carlo iterations per combination = 1,200 runs (evasion parametric sweep). An additional 300 runs from convergence verification (3 sample sizes x 50/100/200 iterations). **Total: 1,500 Monte Carlo runs.**



Results

Finding 1: Ship Destroyed in 100% of Runs

The Burke was destroyed at t=139s in every run across all 1,500 engagements. No interceptor reached terminal phase resolution against the first impacting HGV before ship destruction. Interceptors already in flight continued to resolve posthumously against later-arriving HGVs, producing the posthumous terminal-resolution events discussed below. The engagement timeline is deterministic:

```

t=0s      DF-17 HGV begins modeled terminal segment at ~370 km slant range,
          60 km altitude, Mach 10 initial velocity (3,400 m/s)
          (beyond organic detection range - SPY-6 cannot yet see the target)
t=~110s   HGV enters SPY-6 effective organic detection envelope (~95 km)
          [SPY-6 card: 300 km at 1.0 m^2 reference. Fourth-root RCS scaling:
          300 km x (0.01/1.0)^0.25 = 95 km for 0.01 m^2 HGV]
t=~112s   Track established, threat classified, engagement authority granted
t=~114s   First SM-6 launch command issued
t=~120s   SM-6 at burnout speed (~1,360 m/s), beginning midcourse guidance
t=~133s   SM-6 entering terminal phase - HGV closing at ~4,000+ m/s combined
t=139s    First HGV impacts Burke. Ship destroyed.
          ----- POSTHUMOUS ZONE -----
t=~465s   SM-6 rounds already in flight resolve against remaining HGVs
t=~580s   Last posthumous engagement resolves

```

The total flight time of 139 seconds (versus ~109 seconds for a straight-line slant range at constant Mach 10) reflects two physics effects modeled in the simulation: the curved glide/dive trajectory is longer than the slant range, and atmospheric drag decelerates the HGV significantly during descent through denser air below 30 km altitude. Both effects are defender-favorable - they give the defense slightly more time than a naive straight-line estimate would suggest.

The first DF-17 impact delivers approximately 9.3 GJ of combined kinetic and explosive energy against the Burke's hull structure. The simulation's structural damage model evaluates this as a single-impact kill - the damage from one HGV exceeds the ship's total structural integrity. This is consistent with open-source analysis of hypersonic weapon effects against surface combatants: the Burke was not designed to absorb this class of impact.

All reported intercept rates in this document are posthumous - kills achieved by interceptors that were already in flight at the time the ship was destroyed. These kills are operationally meaningless for the ship that launched them.




Figure 5: Engagement Timeline - Baseline 8-Salvo (Seed 1). Ship destroyed at t=139s. All interceptor kills occur after destruction.

Finding 2: PAC-3 MSE Achieves Zero Kills at All Evasion Levels

Evasion	PAC-3 Dispatches (mean)	PAC-3 Resolutions (mean)	PAC-3 Kills	PAC-3 Hit Rate
0G (ballistic)	12-18	9-13	0	0%
10G	12-18	9-13	0	0%
20G (assessed)	12-18	9-13	0	0%
30G (upper bound)	12-18	9-13	0	0%

Across all variants that carry PAC-3 MSE (Variants B and C), PAC-3 MSE interceptors reached terminal resolution against HGVS targets but achieved zero kills. At 0G - non-maneuvering ballistic reentry, the design point for hit-to-kill interceptors - 450 resolution events across 50 seeds produced zero hits and zero damage registering events.

The failure mechanism is closing velocity, not target maneuver. This is the critical finding.

At Mach 10 threat speed + Mach 5 interceptor speed in head-on geometry, combined closing velocity exceeds 4,000 m/s (approximately 14,400 km/h). Proportional navigation guidance at this closing speed produces miss



distances of 10-50 meters even against perfectly predictable, non-maneuvering targets. The miss distance floor is governed by a fundamental relationship:

Miss distance is proportional to (acceleration authority) / (closing velocity)²

PAC-3 MSE's 40.8G acceleration authority is impressive against Mach 3-4 reentry vehicles where closing velocities are ~2,000 m/s. At that regime, the guidance loop executes 8-12 meaningful correction updates in the final second, driving miss distance below 1 meter - well within the hit-to-kill engagement volume, to say nothing of the Modified Lethality Enhancer.

At Mach 10 closing velocity, the same guidance loop gets 3-5 updates in the final second. Each update corrects position, but the geometry changes faster than the interceptor can steer. The resulting miss distances (10-50m) are an order of magnitude larger than the PAC-3 MSE's 5-meter lethality enhancer radius.

The Modified Lethality Enhancer - 24 titanium cycloid fragments propelled by a 330g shaped charge - was designed as a last-ditch augmentation for near-miss scenarios against theater ballistic missiles. It extends the kill radius from body-to-body contact (~0.3m) to approximately 5 meters. Against hypersonic targets where miss distances start at 10 meters, this is insufficient by a factor of two to ten.

This is not merely an integration limitation. Within the tested geometry, closing velocity dominates the terminal miss-distance budget. Improvements to seeker quality, guidance law, and divert authority can reduce miss distance, but they do not remove the underlying timeline problem: the interceptor must generate lethal geometry before a Mach 10 target reaches the defended asset. A PAC-3 MSE Block II with twice the acceleration authority would narrow the miss distance distribution, but the engagement timeline - not per-round lethality - remains the binding constraint.

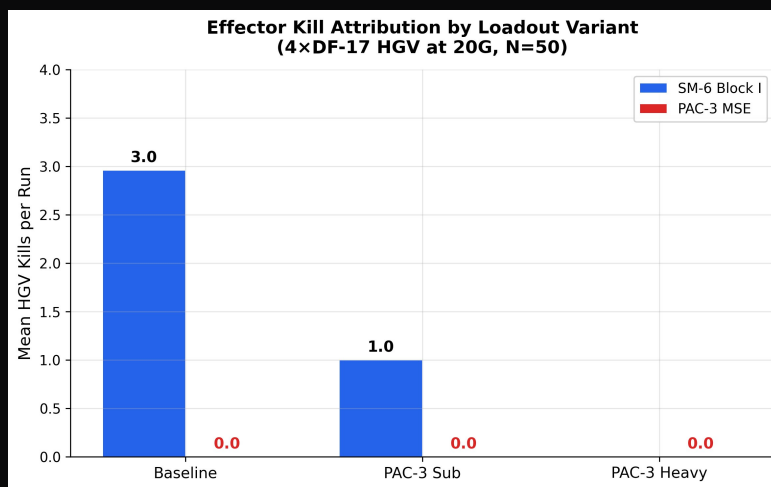


Figure 2: Effector Kill Attribution by Loadout Variant (4x DF-17 HGV at 20G, N=50)



Finding 3: Evasion Paradoxically Improves Defender Performance

Evasion	Baseline IR (4-salvo)	Baseline IR (8-salvo)	SM-6 Kills/Run (mean)	SM-6 Dispatches (mean)
0G (ballistic)	41.5% +/- 11.8%	23.7% +/- 6.8%	1.66	15
10G	71.5% +/- 8.7%	40.9% +/- 5.7%	2.86	21
20G (assessed)	74.0% +/- 4.9%	39.4% +/- 6.1%	2.96	21
30G (upper)	74.0% +/- 4.9%	41.1% +/- 4.6%	2.96	22

The non-maneuvering ballistic target (0G) is harder to kill than the 20G maneuvering target. SM-6 achieves 1.66 kills per run at 0G versus 2.96 kills per run at 20G - a 78% improvement in kills when the target maneuvers.

This was counterintuitive, in a way, it still feels as such. The prevailing narratives around hypersonic defense emphasize that HGV maneuverability is the primary challenge. Our data shows the opposite: speed is the binding variable, and maneuver actually helps the defender. As in most things, *speed kills*.

The mechanism: Random jink maneuver converts a portion of the HGV's straight-line approach velocity into lateral displacement. The HGV still travels at Mach 10 airspeed, but its ground-track closure rate against the ship is reduced. This extends time-in-zone, which is the duration the HGV spends within the SM-6 engagement envelope. More time in zone produces more interceptor dispatch opportunities. At 0G, the fire control loop dispatches 15 SM-6 rounds. At 20G, it dispatches 21 rounds - a 40% increase in engagement attempts.

SM-6's blast-fragmentation warhead has a lethal radius of approximately 25 meters. At the miss distances produced by proportional navigation against a Mach 10 target (even a maneuvering one), the blast-fragmentation radius is large enough to catch the miss distribution. The additional engagement opportunities from extended time-in-zone more than compensate for any increase in per-round miss distance from target jinking.

The effect saturates at 10G. Results at 10G, 20G, and 30G are statistically indistinguishable (71.5% / 74.0% / 74.0% for 4-salvo). Once the HGV is maneuvering at all, the time-in-zone extension dominates. Additional G-loading does not help the attacker because the blast-frag radius is already sufficient to kill at the resulting miss distances.

Implication for threat design: The optimal attacker strategy against blast-fragmentation naval air defense is to not maneuver. A ballistic approach at maximum speed minimizes time in the engagement zone and reduces the number of interceptors that reach terminal resolution. Maneuver is valuable against hit-to-kill interceptors



(where it increases miss distance beyond the kill volume), but counterproductive against blast-fragmentation interceptors (where it increases time-in-zone without exceeding the lethal radius).

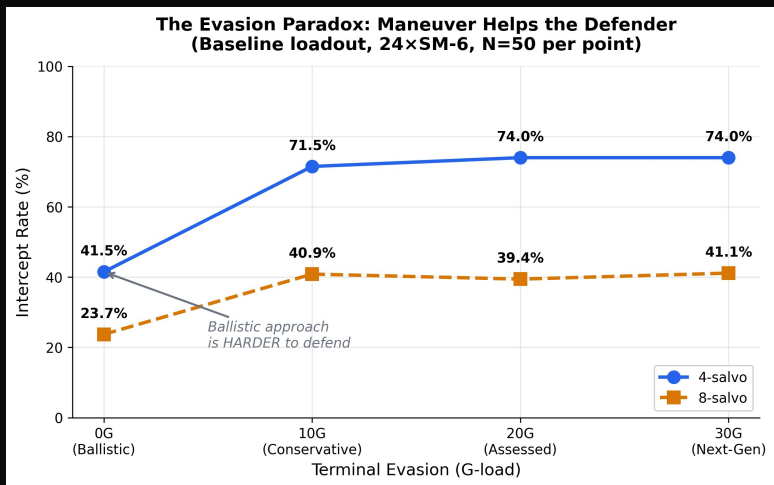


Figure 3: The Evasion Paradox - Maneuver Helps the Defender (Baseline loadout, 24x SM-6, N=50 per point)

Finding 4: Salvo Saturation at 8 HGVs

Results for 8, 12, and 16-HGV salvos are identical within statistical noise for each loadout variant. The Burke is destroyed by the first HGV at t=139s regardless of salvo size. Subsequent HGVs arrive after ship destruction. The simulation continues tracking interceptors already in flight, but no new engagements are initiated post-destruction.

The 4-salvo variant shows higher intercept rates than 8-salvo because fewer total threats means the same number of posthumous kills represents a higher fraction. In absolute terms, SM-6 kills per run converge in such a way that are approximately equal across salvo sizes - the number of interceptors the Burke can put in the air before t=139s is the limiting factor, not the number of incoming threats. As noted earlier, we are engaging at machine speeds likely much faster than a peerless human crew can engage at, in the most favorable conditions.

Finding 5: PAC-3 MSE Integration Degrades Overall Defense

Variant	4-Salvo IR (20G)	8-Salvo IR (20G)	SM-6 Kills/Run	PAC-3 Kills/Run	Mean Cost/Run
A: Baseline (24x SM-6)	74.0%	39.4%	2.96	-	\$108M
B: PAC-3 Sub (12xS M-6+24xPAC-3)	25.0%	16.7%	1.0	0.0	\$63-76M



Variant	4-Salvo IR (20G)	8-Salvo IR (20G)	SM-6 Kills/Run	PAC-3 Kills/Run	Mean Cost/Run
C: PAC-3 Heavy (6x SM-6+36xPAC-3)	0%	0%	0.0	0.0	\$45M

At the assessed DF-17 evasion capability (20G):

Variant B (PAC-3 MSE replacing SM-2) reduces the posthumous intercept rate from 74% to 25% for 4-salvo and from 39% to 17% for 8-salvo. The 24 PAC-3 MSE rounds contribute zero kills. The entire intercept rate reduction comes from halving the SM-6 allocation (24 cells to 12 cells). Every single kill in Variant B is attributable to SM-6.

Variant C (PAC-3 MSE heavy) produces a 0% intercept rate across all evasion levels and salvo sizes. With only 6 SM-6 cells, the Burke dispatches 6 rounds - each individual SM-6 hit delivers approximately 80 units of structural damage against the HGV's 300-unit total integrity, requiring 4 or more rounds to achieve a kill. Six rounds are insufficient to accumulate lethal damage on any single HGV.

The PAC-3 MSE naval integration, as modeled, does not add defensive capability against hypersonic glide vehicles. It displaces capability by consuming VLS cells that would otherwise carry SM-6.

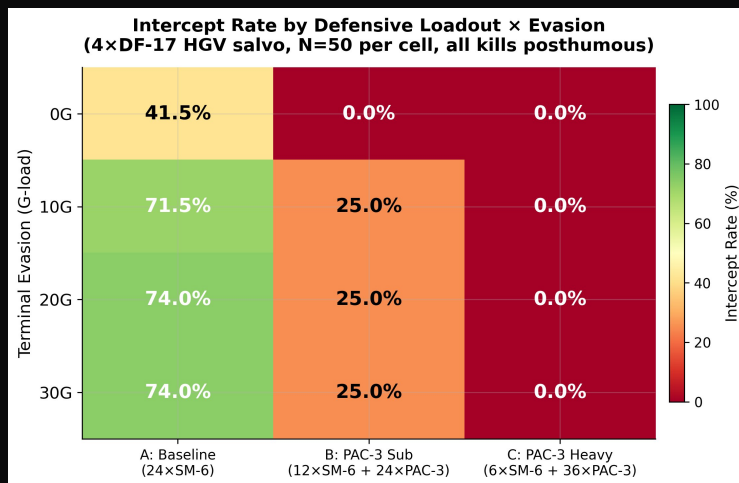


Figure 1: Intercept Rate by Defensive Loadout x Evasion (4x DF-17 HGV salvo, N=50 per cell, all kills posthumous)




Figure 6: Intercept Rate by Variant x Evasion (8-Salvo, N=50 per bar, all kills posthumous)

Cost-Exchange Analysis

Per-Engagement Economics

Variant	Interceptor Expenditure/Run	Ship Replacement	Total Defender Cost	Attacker Salvo (8x DF-17)	Exchange Ratio
A: Baseline	\$130.5M	\$2,200M	\$2,330.5M	\$24M	97:1 attacker
B: PAC-3 Sub	\$76.5M	\$2,200M	\$2,276.5M	\$24M	95:1 attacker
C: PAC-3 Heavy	\$45M	\$2,200M	\$2,245M	\$24M	94:1 attacker

Interceptor costs: SM-6 Block I at \$4.5M per round (MDAA, USN FY2024). PAC-3 MSE at \$4.05M per round (Navy FY2027 budget justification). ESSM Block 2 at \$1.8M per round (not expended against HGV threat class in any run). SM-2 Block IIIA at \$2.1M per round (not expended - cannot engage HGVs).

Threat cost: DF-17 HGV at ~\$3M per round. This is an OSINT estimate based on CSIS Missile Threat Project data and is likely conservative - Chinese weapons procurement costs are opaque, and actual unit costs may be lower at PLARF production volumes.

Ship replacement cost: \$2.2B for a Flight IIA Arleigh Burke DDG based on CBO 2025 shipbuilding cost data. Flight III variants with AN/SPY-6 AMDR run \$2.5-2.7B. The cost of the ship dwarfs the interceptor



expenditure by a factor of 17-49x depending on variant. The interceptor cost is immaterial to the exchange ratio.

At \$3M per DF-17 and \$2.2B per Burke, the attacker achieves a **97:1 cost-exchange ratio** in the baseline case. Eight missiles costing \$24M total destroy a \$2.2B warship plus \$130M in expended interceptors. Even if the DF-17 costs \$10M per round (the upper bound of any credible estimate), the exchange ratio is still 29:1 in the attacker's favor.

The defender cannot win this exchange. The cost-per-kill arithmetic is conclusive: in the best case (Variant A, 4-salvo, 20G), the Burke expends \$108M in SM-6 rounds to achieve 2.96 posthumous kills, a cost-per-kill of \$36.5M against a \$3M target. The kills are posthumous. The ship is destroyed. The \$108M in interceptors protected nothing.

The Exchange Ratio in Context

The 97:1 ratio measures dollars. It does not measure what dollars cannot.

An Arleigh Burke Flight IIA carries a complement of 329. From our understanding, every soul aboard is at a watchstation when the ship goes to general quarters, and most of those stations are below the waterline, behind hatches that are dogged shut - because that is what keeps the ship fighting and what keeps the sea out and what keeps the men in.

At \$3M per DF-17 and 329 lives per Burke, the arithmetic is: **\$72,948 per life**. Less than the cost of a mid-range sedan or used German SUV. That is what the exchange ratio means when you remove the abstraction.




Figure 4: Cost-Exchange Ratio by Loadout Variant (8x DF-17 at 20G, ship + interceptors vs salvo cost)



Observations: PAC-3 MSE Naval Integration Beyond Hypersonic Defense

The headline finding - PAC-3 MSE cannot intercept HGVs - does not mean PAC-3 MSE on a Burke is without value. The FY2027 integration addresses a broad threat spectrum, not exclusively hypersonic glide vehicles. Three properties of the PAC-3 MSE/Aegis combination deserve separate analysis.

Illumination Budget Relief

SM-2 Block IIIA is semi-active radar homing. It requires the AN/SPY-6 to dedicate a continuous-wave illumination beam to the target through terminal phase. In a saturated engagement - eight inbound cruise missiles, for example - the binding constraint on an Arleigh is not magazine depth. It is the number of targets SPY-6 can simultaneously illuminate in terminal. Every SM-2 in terminal flight consumes one of these channels.

SM-6 partially addresses this: its AMRAAM-derived active seeker is fire-and-forget in terminal, consuming only midcourse uplink bandwidth. PAC-3 MSE takes it further. After the initial cue from SPY-6 (S/X-band), the missile's Ka-band active seeker is fully autonomous. From our understanding, this translates to zero illumination load on SPY-6 and zero uplink dependency in terminal phase.

Every PAC-3 MSE in the air is an illumination channel freed for an SM-2 or an additional SM-6 midcourse track. Against the threat classes PAC-3 MSE was designed for - theater ballistic missiles at Mach 3-4, where hit-to-kill guidance achieves sub-meter miss distances - this changes the fire control calculus. A Burke carrying PAC-3 MSE alongside SM-2 can service more concurrent threats than one carrying SM-2 alone, because the PAC-3 MSE engagements do not compete for illumination resources.

Against HGVs, this benefit is moot, PAC-3 MSE cannot achieve a kill regardless. But against the broader threat environment (cruise missiles, tactical ballistic missiles, aircraft), the illumination budget relief is a genuine capability addition.

Electronic Attack Resilience Through Band Diversity

Current Burke air defense is almost entirely S/X-band dependent. SPY-6 operates in S/X-band. SM-2's semi-active homing rides the Aegis illumination. Adversary standoff jamming systems are optimized to degrade S-band and X-band naval radars - because that is what NATO uses to shoot with.



PAC-3 MSE's terminal seeker operates in Ka-band (33-37 GHz). This is a completely different part of the electromagnetic spectrum from SPY-6's S-band (2-4 GHz) and X-Band (8-12 GHz). The jammer burning down SPY-6's detection range does not touch the PAC-3 MSE's terminal seeker. The kill chain gains band diversity: S-band for search, track, and initial cue; Ka-band for terminal kill.

The adversary now must jam two widely separated frequency bands simultaneously to defeat one interceptor. This doubles their EW resource allocation per engagement. SM-2 in a jammed environment degrades twice - once on midcourse guidance (S-band uplink), once on terminal homing (X-band illumination). PAC-3 MSE degrades only on midcourse cue. Terminal guidance is in Ka-band, which current-generation standoff jammers are not built to address.

Anti-Surface Hard-Kill (Speculative)

PAC-3 MSE is a 315 kg missile traveling at Mach 5+ with an active Ka-band seeker that can trivially resolve a ship-sized target. The lethality enhancer sleeve provides blast-fragmentation capability on top of kinetic energy. At 15 km range against a Type 055's bridge or primary radar array, flight time is under 5 seconds. The target cannot maneuver meaningfully in that window.

Nobody designed PAC-3 MSE for anti-surface warfare. But nobody designed SM-6 for anti-surface warfare either, and the AIM-174B air-launched SM-6 variant now exists as a dedicated long-range strike weapon. The precedent is established: missiles designed for air defense can be repurposed for surface attack when the seeker can acquire the target and the kinematics support the engagement.

A Burke that has already expended its Harpoon allocation and is facing a closing surface combatant - a PAC-3 MSE aimed at the bridge superstructure is a Mach 5 kinetic kill weapon with Ka-band terminal guidance. The kinetic energy at impact is in the hundreds of megajoules. That is a mission kill on any superstructure it hits.

This is an author's observation, not a simulation finding. We have not modeled PAC-3 MSE anti-surface engagements. The physics are consistent with the capability; whether doctrine would ever permit it is a separate question.

Discussion

The Engagement Timeline Is the Problem

The central finding of this research is not about interceptor performance, warhead lethality, or seeker technology. It is about time. The DF-17 HGV covers ~370 km of slant range from 60 km altitude in



approximately 139 seconds - longer than the ~109 seconds a straight-line constant-velocity estimate would suggest, because the simulation models atmospheric drag deceleration during descent and the curved glide/dive trajectory. After applying fourth-root RCS scaling to the SPY-6 card value of 300 km against a 1.0 m² reference target, the modeled effective organic detection range against a 0.01 m² HGV is approximately 95 km. At the HGV's descent-degraded velocity, this yields roughly 29 seconds from first organic detection to impact, and approximately 25 seconds after track establishment and fire-control solution latency.

No interceptor in the tested loadout - SM-6 at Mach 3.5, PAC-3 MSE at Mach 5, ESSM at Mach 4 - reaches terminal resolution before this time. The ship is destroyed before the defense can act. This is not a performance gap. It is geometry.

The HGV travels at 3,400 m/s. The fastest tested interceptor (PAC-3 MSE) travels at 1,700 m/s - half the threat's speed. In a head-on engagement, the interceptor must cover the distance to the intercept point before the threat covers the distance to the ship. At a 2:1 speed disadvantage, the interceptor needs to be launched when the threat is at twice the desired intercept range. This requires detection, track establishment, and fire control solution at ranges that challenge even SPY-6 against a 0.01 m² target.

This is not a problem that can be solved by a faster interceptor in the current paradigm. An interceptor traveling at Mach 10 - matching the threat speed - would need to be launched at the moment of detection to achieve a head-on intercept before the threat reaches the ship. With organic detection at ~95 km against a 0.01 m² target, no fire control system operates with zero latency, nor with any interference in a combat scenario against any peer threat.

What the Evasion Paradox Means for Threat Design

The finding that maneuver helps the defender has direct implications for adversary weapon design. If the objective is to defeat blast-fragmentation naval air defense, the optimal terminal profile is maximum speed on a straight line. Maneuver is valuable against hit-to-kill interceptors (where it increases miss distance beyond the kill volume) but counterproductive against blast-fragmentation interceptors (where it increases time-in-zone without exceeding the lethal radius).

A sophisticated adversary - and the PLARF qualifies - would logically develop two terminal profiles: a straight-line ballistic approach for use against SM-6-equipped defenders, and a maneuvering approach for use against hit-to-kill defenders (SM-3, THAAD, future GPI). The DF-17's software-defined terminal guidance likely already supports both modes.

What This Analysis Does Not Address

This study tests one layer of a multi-layer defense architecture. The findings should not be interpreted as "naval defense against HGVs is impossible." They should be interpreted as "terminal kinetic defense by a single ship



against HGVs is insufficient under the tested conditions." The broader defense architecture includes:

Exo-atmospheric intercept. Earlier-layer intercept is not modeled. Exo-atmospheric or high-altitude intercept opportunities - including SM-3 Block IIA engagements where geometry permits - would provide additional defensive layers before the terminal timeline collapses. HGVs specifically complicate exo-atmospheric BMD because they spend much of the relevant flight regime below classic midcourse trajectories and can maneuver to depress their profile. Future research should test layered SM-3 + SM-6 defense with realistic HGV trajectory envelopes.

Cooperative engagement. CEC/NIFC-CA with E-2D Advanced Hawkeye extends detection range to the radar horizon of the Hawkeye's AN/APY-9 radar (significantly further than SPY-6 against low-RCS targets). This may provide 1-2 additional SM-6 engagement opportunities before the terminal timeline collapses. Multi-ship defense with distributed VLS capacity fundamentally changes the engagement calculus.

Directed energy. HELIOS (60 kW class high-energy laser) and future 150+ kW shipboard HEL systems engage at the speed of light. There is no closing velocity problem, no engagement timeline compression, no miss distance governed by guidance bandwidth. The parametric sweep in this study identifies precisely the regime where kinetic intercept faces fundamental physics constraints and directed energy becomes the necessary answer.

Electronic warfare. Not modeled, and its absence uniformly favors the defender. Standoff jamming, onboard ECM, chaff corridors, and decoys would degrade interceptor performance below the levels reported here.

Left-of-launch. Cyber operations against TEL command networks, kinetic strike on launcher positions, space-based persistent tracking of mobile TELs - these address the threat before it enters the terminal engagement zone. The most cost-effective interceptor is the one that destroys the missile on the ground.

Forward Vector

The data points in one direction. Terminal kinetic defense against hypersonic threats - by a single ship, with current-generation interceptors - is a losing proposition. The physics are clear and the exchange ratio is conclusive. The question is not whether this is true. The question is what to do about it.

Speed-of-light engagement. Directed energy weapons eliminate the closing velocity problem entirely. A 150 kW high-energy laser engages at 300,000 km/s. There is no engagement timeline compression, no guidance bandwidth limit, no miss distance. The cost per shot is measured in single-digit dollars. The technology exists. The engineering challenge is beam power, atmospheric propagation, and dwell time against hardened targets -



all solvable problems on a known trajectory. HELIOS is a start. It is not sufficient against ablative-clad, reentry-hardened HGVs, but the scaling path is clear.

Distributed lethality. A single Burke is the wrong unit of analysis. A distributed force of 4-6 surface combatants with CEC connectivity presents 4-6x the VLS capacity, 4-6x the illumination channels, and geometric diversity that forces the attacker to solve a multi-axis targeting problem. The posthumous intercept rates in this study (40-74% per ship) become pre-impact intercept rates when the interceptors are launched from platforms that are not the target.

Left-of-launch. The most favorable exchange ratio is achieved by destroying the DF-17 on its TEL before launch. At \$3M per HGV and \$1.27M per JASSM-ER, the cost exchange inverts to 2.4:1 in the defender's favor - but only if the TEL can be located and targeted before it fires. This is a space-based persistent ISR and kill-chain speed problem, not an interceptor problem.

Autonomous attritable platforms. The human cost analysis in the preceding section identifies the fundamental asymmetry: the attacker risks \$3M in hardware, the defender risks \$2.2B in hardware and 329 lives. Removing the crew from the engagement - through autonomous surface combatants, unmanned distributed VLS barges, or tethered sensor-shooter networks - does not improve the intercept rate, but it changes the exchange ratio by removing the irreplaceable element from the cost column.

None of these are fast. None of them are cheap. But they are the answers that the physics permits, and they are the answers the data in this document demands.

Methodology

Entity Cards

All entity parameters are sourced from open-source intelligence. Cards are version-controlled YAML files with per-field source attribution.

Entity	Card ID	Primary Sources
Arleigh Burke DDG (Flight IIA)	arleigh-burke-ddg	Navy Fact File, CBO shipbuilding data
SM-6 Block I	sm-6-blk1	CSIS Missile Threat, Raytheon published specs



Entity	Card ID	Primary Sources
PAC-3 MSE	pac3-mse-interceptor	CSIS, DTIC ADA319957, TWZ, Navy FY27 budget
DF-17 HGV (20G)	df-17-hgv	CSIS Missile Threat, DoD China Military Power
DF-17 HGV (0G)	df-17-hgv-0g	Research variant
DF-17 HGV (10G)	df-17-hgv-10g	Research variant
DF-17 HGV (30G)	df-17-hgv-30g	Research variant

Convergence Verification

The baseline 8-salvo scenario was run at N=50, 100, and 200 iterations to verify statistical convergence:

N	Intercept Rate	+/- Std Dev	SM-6 Kills/Run
50	39.43%	6.10%	2.76
100	39.57%	6.34%	2.77
200	40.07%	6.01%	2.80

Intercept rate shifted 0.64 percentage points from N=50 to N=200. Standard deviation was stable across all sample sizes. All invariant metrics (dispatches, cost, time-to-first-kill) were identical. N=50 is sufficient for the reported precision level.

Author's Note: The Human Cost Column

There is a column in every cost table that has no header because no one has agreed on the unit of measure. It is not denominated in dollars or rubles or yuan. It does not depreciate. It cannot be amortized. It is the column where the dead go, and it is the column that every exchange-ratio analysis must eventually confront or else confess that it is not serious.

The man who keys the launch sequence for an autonomous one-way attack system - outside of propaganda videos meant to demoralize our warfighters - will never see what he has done. He sits in a climate-controlled compartment or a buried revetment or a truck cab that smells of diesel, tea, and stale tobacco. Emotionless from



the adrenaline dump that came days or hours ago at the start of the engagement, he enters coordinates that are at that moment only numbers and he presses a thing that must be pressed and the weapon leaves and math begins. He will not hear the general quarters alarm on the ship whose coordinates he entered. He will not see the sailors two decks below the waterline who will never reach a ladder. He is, in the antiseptic language of the doctrine, at standoff. He has wagered between twenty thousand dollars in sheet metal and a small piston engine or a multi-million dollar ballistic missile. He did not pay for it. He has wagered no one and nothing.

But not every launch is standoff and not every operator goes home clean.

The crew of a MiG-31K that carries a Kinzhal does not fire from a buried console. They fly the aircraft to the launch point and they press the release and they turn for home and whether they arrive there depends on what is in the air between them and their field. The pilot of a Su-34 delivering FAB-3000 glide bombs over contested airspace has committed his life to the proposition that the suppression of enemy air defenses was thorough and that the intelligence was current and that no one is waiting for him with a Patriot or an AMRAAM. The crew of a Gorshkov-class frigate launching Kalibrs from the Black Sea has bet their ship and two hundred of their countrymen that the Ukrainians cannot reach them. Sometimes they are wrong. The Moskva carried five hundred and ten. Not all of them came home.

The submariner who fires a salvo of cruise missiles from beneath the surface risks less than the surface sailor but he risks everything nonetheless. He lives inside a steel tube that either surfaces or does not. There is no middle outcome. The men aboard a Type 022 Houbei fast attack craft launching YJ-83s in the Taiwan Strait are riding a catamaran the size of a large yacht into the teeth of the most heavily sensed battlespace on Earth. They are young and they volunteered and they know exactly what a Harpoon does to a hull that thin.

This is not sentimentality. It is arithmetic. Every manned platform on every side of every table in our compendium carries a number that does not appear in the cost column. That number is the complement. It is the count of human beings who are aboard when the platform is performing the function for which it was built. The Arleigh Burke carries three hundred and twenty-nine. The Type 055 Renhai carries three hundred and ten. The Admiral Gorshkov carries two hundred and ten. The Nimitz carries five thousand. The Fujian will carry something close to the same. The mathematics of what happens when a hypersonic glide vehicle meets any of these vessels are not abstract. The kinetic energy alone is measured in gigajoules. The warhead is almost beside the point.

The economics of standoff warfare are obscene and alien in their asymmetry and this asymmetry runs in only one direction: toward the unmanned. A Shahed costs twenty thousand dollars and risks no one. The sailor who stands in the path of what it carries has already bet everything he will ever have or be on the readiness of



systems he did not build and the proficiency of men beside him he must trust absolutely because there is no other choice. There has never been another choice. A saturation raid against a carrier does not need to sink her. It needs only to mission-kill the flight deck. Four hits. Maybe three. Five thousand people live under that flight deck.

But the asymmetry is a temporary condition, not a law of physics. The side that transitions its kill chain from manned to autonomous faster does not eliminate the dying. It relocates the dying to the other side of the ledger. This is the strategic logic of directed-energy weapons at three dollars a shot and loitering munitions at fifty thousand dollars a unit and autonomous undersea vehicles that carry no crew and need no rescue. The race is not to build a cheaper missile. The race is to remove the human from the cost column entirely - and to do it before the other side does.

Until that day, every non-autonomous system in this compendium carries two costs. The one you can read is denominated in dollars and will be argued in hearings and water coolers. The one you cannot read is denominated in the only currency that has never inflated and never will. The accountants have tried to capture it. They have assigned values per statistical life. They have built disability-adjusted models and projected lifetime earnings curves and profane actuarials but none of it holds up when you are standing in front of a family explaining what the letter means.

The cost-per-round column does not hold a number for this. There is no number for this.

References

1. Zarchan, P. *Tactical and Strategic Missile Guidance* (7th ed.). AIAA Progress in Astronautics and Aeronautics, 2019.
2. Kingery, C.N. & Bulmash, G. "Air Blast Parameters from TNT Spherical Air Burst and Hemispherical Surface Burst." BRL Technical Report 02555, U.S. Army Ballistic Research Laboratory, 1984.
3. Jeon, D., Kim, K. & Han, S. "Simplified Blast Overpressure Coefficients for Surface Burst." *Computation* 5(3):41, 2017.
4. UFC 3-340-02. "Structures to Resist the Effects of Accidental Explosions." U.S. Department of Defense, 2008.
5. Chandela, R.K. "Fragment Density Distribution from Cylindrical Warheads." *Defence Science Journal* 36(3):331-351, 1986.
6. Sutton, G.P. & Biblarz, O. *Rocket Propulsion Elements* (9th ed.). Wiley, 2017.
7. CSIS Missile Threat Project. <https://missilethreat.csis.org/>
8. "Patriot PAC-3 Missiles to Arm Navy Arleigh Burke-Class Destroyers." *The War Zone*, 2025.
9. Navy FY2027 Budget Justification. Department of the Navy, 2026.
10. DTIC ADA319957. PAC-3 MSE Lethality Enhancement Program. Defense Technical Information Center.



11. Congressional Budget Office. "An Analysis of the Navy's 2025 Shipbuilding Plan." January 2025.
 12. Missile Defense Advocacy Alliance. "Missile Interceptors by Cost." Updated July 2025.
-

Figures

- Figure 1** - Intercept Rate by Defensive Loadout x Evasion G-load (4-salvo heatmap, N=50 per cell)
 - Figure 2** - Effector Kill Attribution: SM-6 vs PAC-3 MSE kills per run by loadout variant at 20G
 - Figure 3** - The Evasion Paradox: intercept rate vs evasion G-load (baseline loadout)
 - Figure 4** - Cost-Exchange Ratio by loadout variant (8-salvo at 20G)
 - Figure 5** - Engagement Timeline: baseline 8-salvo (Seed 1), ship destroyed at t=139s
 - Figure 6** - Variant Comparison: intercept rate across all evasion levels (8-salvo)
-



Appendix: Physics Models - Detailed Derivations

This appendix documents the mathematical models used in the HYPER-DEF-1 simulation, with worked examples specific to the SM-6 vs DF-17 HGV engagement geometry. All formulas are drawn from published, unclassified sources. Implementation-specific calibration, pipeline architecture, and internal model parameters are omitted.

1. Blast Overpressure - Kingery-Bulmash Polynomial

Free-Air Burst (Spherical)

Peak incident overpressure at standoff distance R from a charge of W kg TNT-equivalent:

$$Z = R / W^{(1/3)} \quad [\text{Hopkinson-Cranz scaled distance, m/kg}^{(1/3)}]$$

$$\log_{10}(P_s) = C_0 + C_1*U + C_2*U^2 + C_3*U^3 + C_4*U^4 + C_5*U^5 + C_6*U^6$$

where $U = \log_{10}(Z)$

Coefficients (BRL-TR-02555, 1984):

C0	C1	C2	C3	C4	C5	C6
2.6114	-1.6941	-0.0493	0.2961	-0.0302	-0.0193	0.0031

Valid range: $0.05 \leq Z \leq 40.0 \text{ m/kg}^{1/3}$.

Worked Example: SM-6 Warhead at 15m Miss Distance

SM-6 Block I warhead: 42 kg Comp-B explosive fill. The model uses a Comp-B TNT-equivalence factor of 1.11 for pressure-scaling purposes. TNT-equivalent charge: $42 \times 1.11 = 46.6 \text{ kg TNT}$.

$$Z = 15.0 / 46.6^{(1/3)} = 15.0 / 3.596 = 4.17 \text{ m/kg}^{(1/3)}$$

$$U = \log_{10}(4.17) = 0.620$$

$$\begin{aligned} \log_{10}(P_s) &= 2.6114 + (-1.6941)(0.620) + (-0.0493)(0.620^2) \\ &\quad + (0.2961)(0.620^3) + (-0.0302)(0.620^4) \\ &\quad + (-0.0193)(0.620^5) + (0.0031)(0.620^6) \\ &= 2.6114 - 1.0503 - 0.0190 + 0.0706 - 0.0045 - 0.0029 + 0.0002 \\ &= 1.606 \end{aligned}$$

$$P_s = 10^{1.606} = 40.4 \text{ kPa (free-air)}$$

Surface Burst Enhancement



At low altitude, the detonation produces a hemispherical blast wave with surface reflection. The reflected shock forms a Mach stem that enhances overpressure. Separate empirical coefficients from Jeon, Kim & Han (2017):

$$\begin{aligned}\text{Range 1 (Z = 0.06-1.13): } \log_{10}(P_s) &= 3.133 - 2.171U - 1.388U^2 - 1.040U^3 - 0.371U^4 \\ \text{Range 2 (Z = 1.13-10.0): } \log_{10}(P_s) &= 3.131 - 2.067U - 1.175U^2 + 2.116U^3 - 0.835U^4 \\ \text{Range 3 (Z = 10.0-40.0): } \log_{10}(P_s) &= 3.784 - 4.574U + 2.683U^2 - 0.723U^3\end{aligned}$$

For the SM-6 example (Z = 4.17, Range 2):

$$\begin{aligned}\log_{10}(P_s) &= 3.131 - 2.067(0.620) - 1.175(0.620^2) + 2.116(0.620^3) - 0.835(0.620^4) \\ &= 3.131 - 1.282 - 0.452 + 0.504 - 0.123 \\ &= 1.778\end{aligned}$$

$$P_{s_surface} = 10^{1.778} = 59.9 \text{ kPa}$$

Enhancement factor at this geometry: $59.9 / 40.4 = 1.48x$. This is within the documented 1.3-3.8x range for surface bursts (UFC 3-340-02).

Altitude-Dependent Blending

The simulation blends between surface burst and free-air models based on detonation altitude, transitioning smoothly across the relevant altitude band. For HYPER-DEF-1, all engagements occur at altitude (HGVs at 30-60 km descending, interceptors at lofted altitude). Surface burst enhancement does not apply. The free-air model is used exclusively.

2. Blast Damage Thresholds - Material Basis

DF-17 HGV Resistance

The HGV blast resistance thresholds are calibrated to place the target above ordinary aircraft blast vulnerability while still allowing catastrophic failure when fragment perforation enables internal blast coupling. Carbon-carbon composite panels begin delamination at ~700-900 kPa static overpressure. The TPS must survive high dynamic pressure during terminal descent - at Mach 10, dynamic pressure is roughly ~100 kPa at 30 km altitude and ~500 kPa near 20 km altitude under standard-atmosphere assumptions.

Arleigh Burke DDG Resistance

Onset threshold: Steel hull plating (HY-80/HY-100 steel, 12-25mm thickness) begins permanent deformation at ~50-70 kPa (UFC 3-340-02 Table 5-12). Destruction threshold: Catastrophic hull breach at overpressure far exceeding the threshold at contact range. The ship is destroyed by a single hit regardless of hull location.

Confined Blast Enhancement

When fragments perforate the target before the blast wave arrives (fragment velocity ~2,600 m/s >> blast wave velocity ~340 m/s at 15m), the blast wave enters through perforations and reflects internally. The confined blast



amplification factor is derived from UFC 3-340-02 for partially vented structures.

3. Fragmentation - Gurney Velocity and Lethality Envelope

Fragment Initial Velocity (Gurney Equation, Cylindrical)

$$V_f = \sqrt{2E} \times \sqrt{\left(\frac{C}{M} \right) / \left(1 + 0.5 \times \frac{C}{M} \right)}$$

where:

$\sqrt{2E}$ = Gurney constant (m/s): Comp-B = 2,710 m/s

C = explosive fill mass (kg)

M = casing mass (kg)

SM-6 warhead: C = 42 kg Comp-B, M = 22 kg casing.

$$C/M = 42/22 = 1.909$$

$$\begin{aligned} V_f &= 2710 \times \sqrt{1.909 / (1 + 0.5 \times 1.909)} \\ &= 2710 \times \sqrt{1.909 / 1.955} \\ &= 2710 \times \sqrt{0.977} \\ &= 2710 \times 0.988 \\ &= 2,677 \text{ m/s} \end{aligned}$$

Fragment Velocity at Range (Drag Deceleration)

$$V_f(r) = V_{f0} \times \exp(-C_d \times \rho_{\text{air}} \times A_f \times r / (2 \times m_f))$$

where:

C_d = 0.47 (sphere drag coefficient for natural fragments)

ρ_{air} = 1.225 kg/m³ (sea level)

A_f = presented area of characteristic fragment

m_f = characteristic fragment mass

At 15m miss distance, fragment kinetic energy remains well above the perforation threshold for carbon-carbon composite and ablative TPS materials, enabling the confined blast enhancement described above.

4. Seeker Physics - Acquisition and Track

The simulation models seeker acquisition as a deterministic geometric check based on range-to-target, angle off boresight relative to the seeker field of view, and environmental factors. Seeker parameters are derived from open-source specifications for each interceptor type.

Terminal Pulse Motor (PAC-3 MSE)

PAC-3 MSE has a terminal pulse motor that adds 200 m/s delta-V at terminal phase entry. Relative to a ~4,400 m/s closing velocity, the 200 m/s pulse changes closure by only ~4.5%. The pulse is physically meaningful for reducing miss distance in subsonic/supersonic engagements but insufficient to overcome the fundamental guidance bandwidth limitation at hypersonic closing speeds.



5. Detection Physics - Radar Equation

Fourth-Root RCS Scaling

$$R_{\text{detect}} = R_{\text{ref}} \times (\sigma_{\text{target}} / \sigma_{\text{ref}})^{(1/4)}$$

SPY-6(V)1 parameters: $R_{\text{ref}} = 300,000$ m, $\sigma_{\text{ref}} = 1.0$ m². DF-17 HGV: $\sigma_{\text{target}} = 0.01$ m² (low-observable glide body).

$$\begin{aligned} R_{\text{detect}} &= 300,000 \times (0.01 / 1.0)^{(1/4)} \\ &= 300,000 \times 0.316 \\ &= 94,870 \text{ m, approximately } 95 \text{ km} \end{aligned}$$

At 95 km detection range, the HGV has already been in flight for approximately 110 seconds (having covered ~275 km of its ~370 km slant-range trajectory). With impact at $t=139$ s, this leaves approximately 29 seconds from first organic detection to impact. Subtracting track establishment (~2s) and fire control solution (~2s), the engagement window is approximately 25 seconds.

Atmospheric Attenuation (ITU-R P.676 + P.838)

Clear-air attenuation: ~0.007 dB/km (ITU-R P.676, 75% RH)

Rain attenuation at 0 mm/hr: 0 dB/km

$$L_{\text{total}} = 0.007 \times 95 \times 2 = 1.33 \text{ dB (two-way)}$$

Atmospheric attenuation reduces effective detection range by ~7% at S-band in clear conditions. After attenuation, effective detection range is ~88 km, corresponding to approximately 26 seconds before impact. With 4 seconds for track establishment and fire control solution, the effective engagement window is approximately 22 seconds.

6. Sea Clutter - NRL Improved Model

$$\begin{aligned} \sigma_0(\text{dB}) &= C1 + C2 \cdot \log_{10}(\sin(\psi)) \\ &+ (27.5 + C3 \cdot \psi) \cdot \log_{10}(f) / (1 + 0.95 \cdot \psi) \\ &+ C4 \cdot (SS+1)^{1/(2 + 0.085 \cdot \psi + 0.033 \cdot SS)} \\ &+ C5 \cdot \psi^2 \end{aligned}$$

where:

ψ = grazing angle

f = radar frequency (GHz)

SS = sea state (0-6)

Coefficients (HH polarization): $C1=-73.0$, $C2=20.781$, $C3=7.351$, $C4=25.65$, $C5=0.00540$. Valid: 0.5-35 GHz, 0.1-60 degrees grazing, SS 0-6. Source: IEEE Trans. AES 48(4), 2012.

7. Proportional Navigation - Guidance Law



$$a_{cmd} = N \times V_c \times (\omega \times r_{hat})$$

where:

N = navigation constant (3 midcourse, 5 terminal)

V_c = closing velocity (m/s)

ω = LOS rate vector = $(r \times v_{rel}) / |r|^2$

r_hat = unit LOS vector

Miss Distance Drivers at Hypersonic Closing Velocity

$$miss_{residual} \approx 0.5 \times a_{target} \times \tau_{lag}^2$$

Effective lag tau_lag	Residual displacement
0.10 s	0.98 m
0.25 s	6.1 m
0.50 s	24.5 m
0.70 s	48.0 m

This supports residual terminal miss distances in the 10-50 m range when the target performs high-G lateral jinks during the final intercept interval and the interceptor is control-authority limited.

8. Energy Management - Interceptor Flight Profile

Tsiolkovsky Boost Phase (Analytical)

$$V_{boost} = V_{launch} + a_0 \times T_{burn} \times (-\ln(1 - f)) / f$$

where:

V_launch = cold-launch ejection velocity (SM-6: 30 m/s from Mk 41 VLS)

a_0 = boost acceleration

T_burn = boost duration

f = propellant mass fraction

The analytical boost velocity is capped to the interceptor's structural and aerothermal speed limit. The SM-6 cannot sustain speeds above approximately Mach 4 due to radome and control surface heating constraints. After burnout, the interceptor coasts with Mach-dependent aerodynamic drag using ISA atmospheric density.

9. Warhead Impact - DF-17 HGV vs Burke DDG

Impact Energy

DF-17 HGV: 500 kg warhead, 315 kg Comp-B, Mach 10 (3,400 m/s) at impact.

$$Kinetic\ energy: 1/2 \times 500 \times 3400^2 = 2.89 \times 10^9\ J = 2.89\ GJ$$



315 kg Comp-B chemical energy: $315 \times 1.11 \times 4.184 \text{ MJ/kg} = 1,462 \text{ MJ} = 1.46 \text{ GJ}$

Total energy delivery: ~4.35 GJ

The kinetic energy exceeds the chemical energy by 2:1.

Overpressure at Contact

$W_{TNT} = 315 \times 1.11 = 349.65 \text{ kg}$

$Z = 0.5 / 349.65^{(1/3)} = 0.5 / 7.04 = 0.071 \text{ m/kg}^{(1/3)}$

$U = \log_{10}(0.071) = -1.149$

$\log_{10}(Ps) \sim 4.037$

$Ps \sim 10,900 \text{ kPa}$

This is near the lower validity bound of the Kingery-Bulmash polynomial ($Z_{min} = 0.05$), so the exact overpressure should not be over-interpreted. The model only needs to establish that the Burke's destruction threshold is exceeded by a large margin.

Summary of Model Fidelity

Model	Basis	Fidelity	Limitation
Blast overpressure	Kingery-Bulmash (BRL-TR-02555)	High	Analytical, not hydrocode
Surface reflection	Jeon et al. 2017	High	Piecewise polynomial fit
Fragmentation	Gurney velocity + calibrated envelope	Medium	Simplified fragment dist.
Blast thresholds	Material properties (UFC 3-340-02)	Medium	Simplified threshold model
PN guidance	Zarchan 3D TPN	High	Perfect info assumed
Energy mgmt.	Tsiolkovsky + ISA + Mach-dependent drag	High	No thermal constraints
Seeker	Geometric acquisition	Medium	No track jitter modeled
Evasion	Random jink with hold intervals	Medium	No optimal evasion strategy
Detection	Fourth-root radar eq. + NRL clutter	High	No terrain masking
Atmospheric	ITU-R P.676/P.838	High	Simplified coefficients



This analysis uses physics-grounded simulation models with OSINT-derived parameters. It is not an operational assessment. Results represent upper-bound performance estimates under defender-favorable assumptions. The authors have no affiliation with any publication, defense contractor nor any government agency referenced herein.

UNCLASSIFIED // OPEN SOURCE INTELLIGENCE ONLY



EMPYREAN DEFENSE

Decision Dominance Through Sensor Fusion

<https://empyreandefense.com/>

This document is UNCLASSIFIED and contains no controlled or restricted information. All data sources are open-source and publicly available. Distribution is unlimited.

UNCLASSIFIED // OPEN SOURCE INTELLIGENCE ONLY