FARMER LABS / ARGUS ENGINEERING STUDY

ARGUS-PD/POD-7K Launcher Science and Application Envelope

Distributed Pulse-Module Architecture, Counter-Impulse Closure, Projectile-Family Physics and Space-Only Research Utility

Christopher I. V. Farmer

Standalone ARGUS engineering study • Independent preprint

JULY 2026

POD-7K functional architecture - duties, not selected hardware

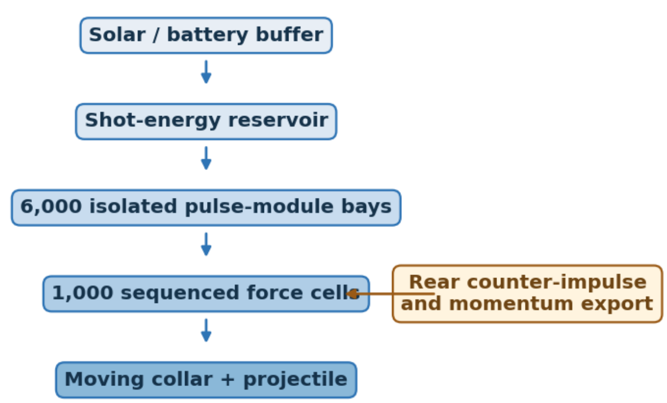

Metrology, quench protection, containment and energy recovery span every tier.

projectile physics • 6,000 module bays • flight environment • scientific utility

Abstract

This standalone engineering study asks what a seven-kilometre electromagnetic launcher and its projectile family would be scientifically useful for if the mechanism could be made to work. Asteroid deflection is retained as one worked application, not as the organising purpose of the paper. The primary contribution is a falsifiable operating envelope linking complete moving mass, release speed, kinetic energy, momentum, specific target energy, pulse-module duty and platform recoil.

Two different control spaces are separated. First, a 150 GJ iso-energy family spans 0.300 kg at 1,000 km s^{-1} to 30 kg at 100 km s^{-1}. Every member carries the same kinetic energy, but direct momentum rises tenfold as speed falls. R-DART therefore favours extreme impact-speed and shock-regime experiments; T-DART favours coherent impulse transfer; V-DART is an intermediate comparison article. Second, a variable-energy sequence for the 0.300 kg article spans 0.15, 1.5, 15 and 150 GJ. Release speed ranges from 31.6 to 1,000 km s^{-1}, while mean admitted power per one of 6,000 equal accounting modules ranges from 0.161 MW to 5.102 GW. The lower modes are still well beyond many conventional test regimes, but they provide a staged path for impact diagnostics, target-response studies and pulse-chain qualification.

The proposed fixed structure is expressed as 1,000 seven-metre force cells served by 6,000 fault-contained pulse-module bays. Six thousand physical enclosures, interfaces and isolation domains are a plausible architectural count; six thousand existing 85.714 MJ pulse modules are not shown to exist. A Tesla Megapack comparison supports a slow battery-buffer tier, not the firing source. Approximately 37 current four-hour units contain the full 514.286 GJ nameplate energy, but their combined published rating is only about 36 MW against a 30.612 TW mean admitted R-DART shot. Present millisecond capacitor cans close the waveform category but imply prohibitive mass when naively scaled.

The rear-valve proposal is valid only as momentum export. A 3 MN s T-DART shot gives a 2,000 t free platform an uncompensated velocity change of 1.5 m s^{-1} and about 0.105 m of constant-force translation during the 140 ms launch. An exactly opposed exhaust impulse can eliminate final centre-of-mass drift in the ideal model, but requires 666.7 kg expelled at 4.5 km s^{-1} and does not eliminate structural reaction.

Space is not a complete vacuum. Interplanetary dust, meteoroids, residual gas and near-planet particulates become impactors in the projectile frame. At 1,000 km s^{-1} a 1 µg grain carries 500 J of local collision energy and the kinetic energy per atomic mass unit is about 5.18 keV. Published hypervelocity research supports cratering, vapour and impact-plasma formation, but no experiment validates the survival of a macroscopic boron-doped-diamond article at 100-1,000 km s^{-1}. A dust contact does not automatically release all 150 GJ; it can create a local damage and plasma source capable of initiating fracture. A CVD-diamond linear-elastic fracture-energy calculation is added as a deliberately conservative lower-bound screen, together with a Poisson encounter model over distance. Neither is a ballistic limit. Catastrophic and cumulative-damage probability must be computed from trajectory-specific flux, grain-size distribution, material flaws and shielding or sacrificial geometry. Diamond, boron-doped diamond, diamond-SiC, toughened SiC, boron

carbide, tungsten and carbon-rich sacrificial layers are compared rather than assuming that one material is universally optimal.

The strongest independent scientific cases are controlled hypervelocity impact and shock research, comparative momentum-coupling and fragmentation experiments, and spacecraft-material studies at purpose-built subscale. Small-body intervention, airless-body geophysics and momentum-beam concepts remain conditional applications. Orbital-debris impact is not treated as a benefit because fragmentation can worsen the hazard. Terrestrial, active satellite and coercive applications are excluded. No buildable launcher, validated full-energy projectile or operational mission is claimed.

Keywords: electromagnetic launcher; pulse power; hypervelocity impact; shock physics; equation of state; momentum coupling; projectile design; interplanetary dust; atmospheric ablation; boron-doped diamond; counter-impulse; planetary science; planetary defence; feasibility.

1. Contribution, scope and civil boundary

The companion ARGUS-PD/POD-7K sequence corrects the moving-mass ledger, defines the integrated seven-kilometre architecture and places the concept on an evidence ladder [1-3]. This paper is separate from that sequence. It brings together four questions that deserve a distinct scientific treatment:

1. What physical quantities can the launcher vary, and which experimental questions does each quantity answer?
2. Which projectile is preferable when the objective is impact speed, delivered energy, direct momentum or target response?
3. Can 6,000 become a defensible hardware architecture, and can a rear valve control the platform's retained momentum?
4. Which proposed uses are supported by analogue evidence, which remain hypotheses, and which should be excluded?

Hypervelocity launchers already serve recognised scientific functions. NASA's Hypervelocity Impact Technology facility reports thousands of tests in impact physics, materials and structural response [4]. NASA has separately described impact testing as a materials-science tool [5], and ESA uses hypervelocity experiments to study shielding and debris-impact protection [6]. Earlier electromagnetic-launcher programmes also identified high-pressure equation-of-state measurements as a principal use [7]. These precedents do not validate POD-7K: conventional facilities generally operate many orders of magnitude below 150 GJ and at far lower macroscopic projectile speed. They do, however, establish that a controlled impact source can be a scientific instrument rather than only a mission system.

This paper deliberately withholds operational detail. It contains no target coordinates, active-satellite cases, firing procedures, command logic or trajectory optimiser. Applications are restricted to civil research, inert or authorised scientific targets, and generic small-body screening. Any future experiment would require facility safety, debris containment or clearance, planetary protection where relevant, independent dual-use review and lawful authorisation.

2. Physics of utility

2.1 Energy, momentum, force and target scale

For complete moving mass m, release speed v, active length L, electrical-to-kinetic efficiency η_e, installed reserve r and equal module count N_m,

$$E_k=\frac{1}{2}mv^2,\ p=mv,\ \acute{F}=\frac{E_k}{L},\ t_L=\frac{2L}{v}. \quad (1)$$

$$E_{admit}=\frac{E_k}{\eta_e},\ E_{nameplate}=(1+r)E_{admit}. \quad (2)$$

$$E_{mod,nameplate}=\frac{E_{nameplate}}{N_m},\ \acute{P}_{mod}=\frac{E_{admit}}{N_m t_L}. \quad (3)$$

The controlled launcher values are $L=7,000$ m, $\eta_e=0.35$, $r=0.20$ and $N_m=6,000$. These are study assumptions, not measured performance.

The usefulness of an impact cannot be inferred from release speed alone. For a target or instrumented specimen of mass M_t, two first-order measures are

$$Q=\frac{E_k}{M_t},\ \Delta v_t=\frac{\beta p}{M_t}, \quad (4)$$

where Q is specific delivered kinetic energy and β is an experimentally determined momentum-coupling factor. Equation (4) is only an accounting boundary. Energy partition into shock, heat, phase change, ejecta, fragmentation and projectile residue depends on relative speed, incidence, geometry and both materials. A meaningful experiment therefore declares its desired observable before selecting a projectile.

2.2 Full-energy projectile family

The five reference articles lie on a common 150 GJ kinetic-energy contour.

Case	Complete mass	Speed	Momentum	Transit	Mean admitted power per module	Primary comparative value
R-DART	0.300 kg	1,000 km s^{-1}	0.300 MN s	14 ms	5.102 GW	highest impact speed; lowest recoil
Q-DART	1.200 kg	500 km s^{-1}	0.600 MN s	28 ms	2.551 GW	high-speed

Case	Complete mass	Speed	Momentum	Transit	Mean admitted power per module	Primary comparative value
						intermediate
V-DART	3.333 kg	300 km s^{-1}	1.000 MN s	46.7 ms	1.531 GW	balanced comparison point
M-DART	13.333 kg	150 km s^{-1}	2.000 MN s	93.3 ms	0.765 GW	momentum-rich intermediate
T-DART	30.000 kg	100 km s^{-1}	3.000 MN s	140 ms	0.510 GW	maximum direct impulse

All five require 428.571 GJ expected admitted energy and 514.286 GJ installed nameplate under the common assumptions. Their scientific roles nevertheless differ:

- **R-DART** is the limiting speed article. If target containment and diagnostics could be solved, it would probe an extreme shock regime while imposing the lowest direct platform impulse. It is the most demanding full-energy pulse-power case.
- **V-DART** provides an intermediate state in both speed and momentum. A comparative R/V/T campaign would be more scientifically useful than a single spectacular velocity point because it separates speed-driven from impulse-driven response.
- **T-DART** maximises direct momentum and reduces instantaneous module power. It is the preferred family member for a coherent Δv objective when arrival time, moving-article mass and recoil are acceptable.

At fixed energy,

$$p=\frac{2E_k}{v}. \tag{5}$$

The equation explains why “most effective projectile” has no universal answer. The R-DART extreme is best when impact speed or shock state is the independent variable. The T-DART extreme is best when direct impulse is the independent variable. Equal energy does not imply equal crater, ejecta, fragmentation, radiation or momentum enhancement.

2.3 Variable-energy operating sequence

Scientific qualification should not begin at 150 GJ. For the same 0.300 kg complete moving article, let f_E be a fraction of the 150 GJ reference energy. Then

$$v(f_E)=v_{full}\sqrt{f_E},\acute{F}(f_E)=f_E\acute{F}_{full}, \tag{6}$$

$$t_L(f_E)=\frac{t_{L,full}}{\sqrt{f_E}},\acute{P}_{mod}(f_E)=f_E^{3/2}\acute{P}_{mod,full}. \tag{7}$$

Mechanical fraction	Mechanical energy	Speed	Momentum	Transit	Installed energy/module	Mean admitted power/module
0.1%	0.150 GJ	31.6 km s^{-1}	0.00949 MN s	442.7 ms	0.0857 MJ	0.161 MW
1%	1.50 GJ	100 km s^{-1}	0.0300 MN s	140.0 ms	0.857 MJ	5.102 MW
10%	15.0 GJ	316 km s^{-1}	0.0949 MN s	44.3 ms	8.571 MJ	161.3 MW
100%	150 GJ	1,000 km s^{-1}	0.300 MN s	14.0 ms	85.714 MJ	5,102 MW

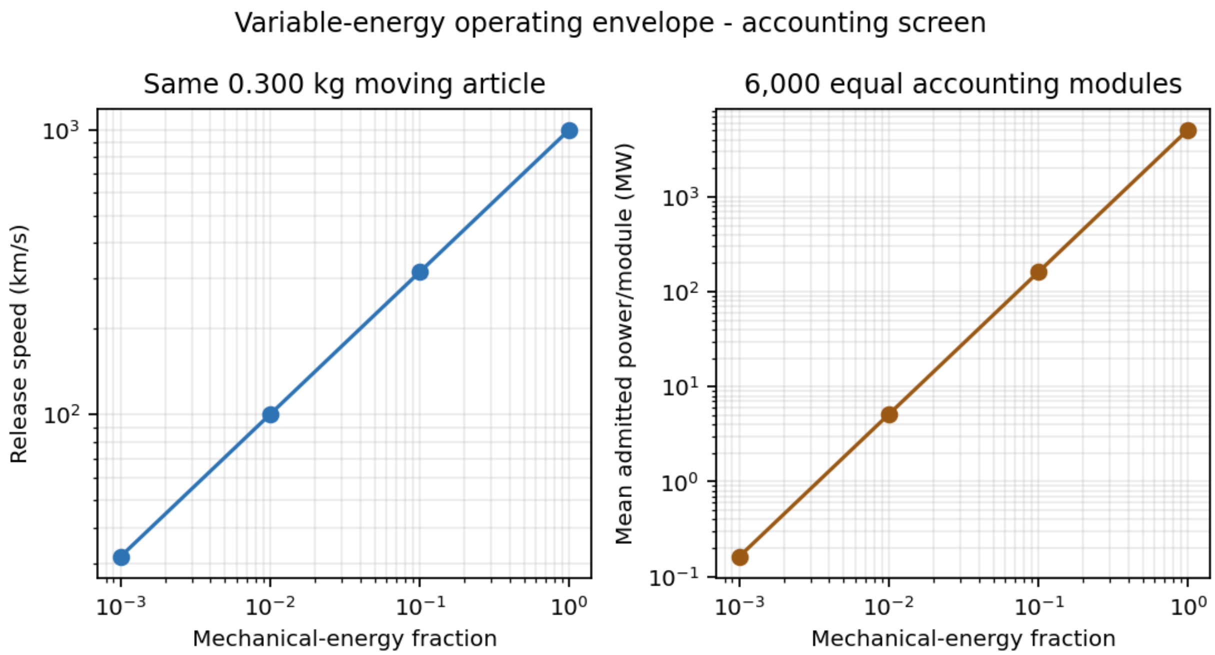

Variable-energy operating modes for a 0.300 kg complete moving article. The curves are accounting relationships, not demonstrated launcher performance.

Even the 0.1% point is a 150 MJ macroscopic impact and cannot be treated as a routine laboratory shot. The table is valuable because it exposes a staged logic: prove synchronisation and metrology at still lower laboratory scale, then progress through declared energy fractions only when target containment, diagnostics, energy transfer and structural response are already measured. The same efficiency cannot simply be assumed at every fraction.

3. Flight environment and projectile state

3.1 Bulk speed is not material temperature

A body moving uniformly through an ideal vacuum does not become hot or weak merely because its centre-of-mass speed is high. In its co-moving frame the boron-doped-diamond article retains a material state set by temperature, stress, defects, doping, charge and radiation history. At 1,000 km s^{-1} the special-relativistic correction to classical kinetic energy is approximately 5.6×10^{-6}, so classical mechanics remains adequate for the launcher ledger.

The threats arise from how that speed is created and from matter encountered after release:

- axial and transverse stress waves during acceleration;
- eddy-current, dielectric or contact heating at the moving interface;
- residual-gas and plasma interaction near the launcher;
- interplanetary or orbital dust impacts;
- charge and ionised impact-plume effects; and
- catastrophic atmospheric energy deposition if a trajectory intersects a planetary atmosphere.

The distinction matters. “Velocity changes the diamond’s properties” is not the correct constitutive statement. Acceleration load and collision-induced internal energy change them.

3.2 What is known about boron-doped diamond

Boron doping is introduced to alter diamond’s electrical properties, but boron concentration and defect structure also change the mechanical and thermal properties that matter to a projectile. Liu et al. measured single-crystal boron-doped diamond and found that material below about 300 ppm boron was close to natural diamond in elastic response, whereas 2,000-3,000 ppm material showed reductions of roughly 2.4%-3.3% in the reported single-crystal elastic constants and a much larger hardness reduction [20]. Prikhodko et al. measured diamond containing about 10^{19} cm^{-3} boron and found thermal conductivity about tenfold below high-quality type-IIa diamond near 100 K, with a difference below 30% above room temperature [21].

These results reject the use of “boron-doped diamond” as a single material property set. A credible moving-article specification must declare at least:

- boron concentration and spatial uniformity;
- single-crystal, polycrystalline or coating architecture;
- crystallographic orientation;
- grain size, inclusions and flaw population;
- electrical resistivity over the firing temperature range;
- thermal conductivity and heat capacity over the same range;
- tensile, compressive and fracture-toughness allowables;
- residual stress after machining, bonding and doping; and
- radiation, vacuum and thermal-cycle history.

The reference core mass is 0.150 kg. If it is idealised as a uniform 50 mm-vertex-diameter diamond article, its calculated length is about 21.9 mm. A one-dimensional end-loaded stress scale is

$$\sigma_{ax,screen}=\rho_c a L_c. \qquad (8)$$

Using $\rho_c = 3{,}510$ kg m^{-3}, $L_c = 21.9$ mm and the mean seven-kilometre acceleration gives approximately 0.0055, 0.0549, 0.549 and 5.49 GPa for the 31.6, 100, 316 and 1,000 km s^{-1} modes. This is not a strength prediction. Distributed collar loading could reduce the end-loaded scale, while pulse asymmetry, stress-wave superposition and microscopic flaws could make local tension worse. The full-speed article requires transient anisotropic fracture analysis and a statistically characterised batch of the specified boron-doped material.

3.3 Candidate-material down-selection

No measured property supports treating boron-doped diamond as the automatic structural winner. The material decision must separate at least four duties: the acceleration load path, the electromagnetic or electrical function, local dust-impact tolerance and the intended target-coupling state.

Candidate	Evidence that supports study	Evidence that prevents selection today	Best present role
CVD or polycrystalline diamond	exceptional stiffness and thermal transport; measured CVD toughness about 5-6 MPa m□$^{1/2}$ [26]	brittle flaw sensitivity; thin-film studies show strong microstructure dependence; no 100-1,000 km s^{-1} ballistic limit [27]	structural coupon and thermal-spreader candidate
Boron-doped diamond	controllable electrical function; low-boron material can remain near undoped elastic response [20,21]	high boron concentration reduces elastic/hardness response; dynamic fracture and dust allowables absent	thin or graded functional layer
Diamond-SiC	10.4 ± 4.0 GPa microscale interface-fracture strength reported for one reactively infiltrated material [28]	scatter, residual silicon, scale and whole-article high-rate response unqualified	scalable hybrid-core candidate
Toughened SiC	ABC-SiC laboratory specimens reached 6.2-7.2 MPa m□$^{1/2}$ room-temperature toughness, well above a commercial SiC comparator [29]	grain-boundary processing controls behaviour; static/fatigue values are not extreme-impact allowables	structural ceramic and load-spreader comparator
Boron carbide	low density, hardness and thermal stability	impact-induced local amorphisation can initiate cracking [30]	low-mass hard-ceramic comparator
Tungsten or tungsten alloy	dense, compact and supported by high-temperature and fracture data	severe moving-mass penalty; room-temperature fracture can be brittle and microstructure-dependent [31]	metal core, collar or sacrificial-nose comparator
Carbon-carbon or graphite-rich layer	low atomic number and ultrafast-probe literature suggesting lower heavy-ion track damage in graphite than quartz [33]	anisotropy, porosity, joining and erosion; interstellar calculations are not an ARGUS ballistic limit	replaceable sacrificial nose or erosion layer
Graded multi-material article	separates load, conduction and shielding functions	residual stress and interfaces may become the dominant failure sites	architecture hypothesis requiring coupon tests

The table changes the selection question from "is diamond strongest?" to "which architecture retains the required state after acceleration and a declared particulate fluence?" At fixed complete mass and speed, material choice does not change total kinetic energy or momentum. It changes projected area, stress-wave impedance, flaw sensitivity, phase response, erosion, conductivity and target contamination. The down-selection therefore requires matched geometries, declared dopant and microstructure, measured high-rate constitutive data, dust-impact tests and post-launch inspection. Static hardness alone is not a selection metric.

3.4 Space dust, meteoroids and residual gas

NASA's Meteoroid Engineering Model 3 resolves flux by direction, speed, mass and density for spacecraft engineering [17]. That class of model is required because "space dust" is not a uniform fog. Risk depends on heliocentric trajectory, pointing, flight time, exposed area, local environment and the critical grain size for the article.

Two adjacent research programmes are especially relevant. Iyer, Mehoke and Batra developed a data-anchored shock-hydrocode ballistic-limit method for a specific Solar Probe Plus two-wall shield over 7-150 km s^{-1} [32]. Its configuration-specific incipient-spall diameters were hundreds of micrometres; those numbers cannot be transferred to an unshielded 50 mm projectile, but the combined test, equation-of-state, hydrocode and uncertainty method can. At much higher speed, Hoang et al. modelled dust and gas damage to a relativistic wafer and predicted cratering, melting and cumulative erosion [33]. That interstellar column and flat-wafer geometry are also non-transferable, but the work confirms that both rare large grains and accumulated small-particle exposure must be treated.

For a grain of mass m_d and relative velocity v_r, the local collision energy scale is

$$E_d = \frac{1}{2} m_d v_r^2. \tag{9}$$

Grain mass	Approximate diameter at 2,500 kg m^{-3}	Collision energy at 100 km s^{-1}	Collision energy at 1,000 km s^{-1}
1 pg	0.913 µm	5 µJ	0.5 mJ
1 ng	9.14 µm	5 mJ	0.5 J
1 µg	91.4 µm	5 J	500 J
1 mg	914 µm	5 kJ	500 kJ

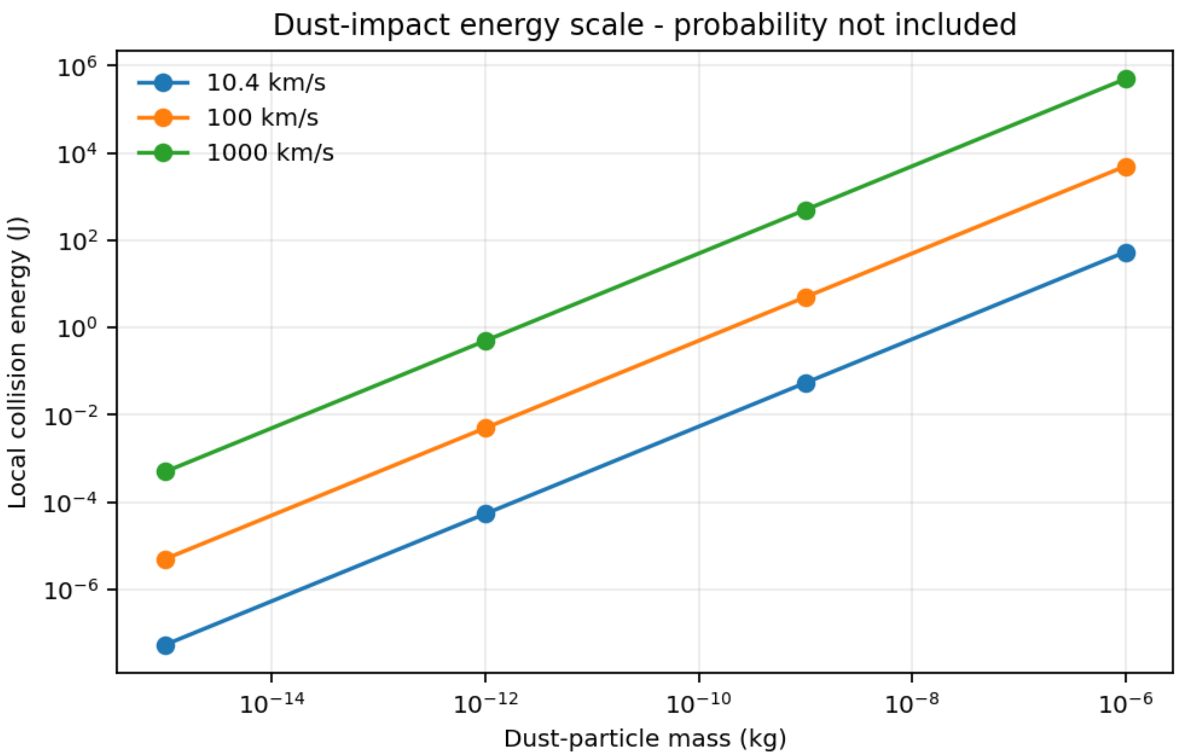

Local dust-collision energy as a function of grain mass and relative speed. The figure contains no impact probability.

NASA research describes hypervelocity impacts as producing vaporisation and ionisation of both impactor and target material [18]; laboratory dust-impact work also observes craters together with vapour and plasma plumes [19]. At 1,000 km s^{-1} the collision energy per atomic mass unit is about 5.18 keV, compared with 51.8 eV at 100 km s^{-1} and 0.56 eV at the demonstrated 10.4 km s^{-1} macroscopic-launcher analogue. The highest ARGUS states are therefore outside the validated solid-impact regime and must be treated as high-energy-density, ionising interactions.

One grain does not automatically disperse all projectile energy. It deposits energy over a small region, creates a crater and expanding ionised plume, and launches stress waves. Total failure occurs only if that disturbance exceeds a flaw-dependent fracture or spall threshold, removes critical geometry, or initiates an unstable breakup. Multiple smaller impacts can erode or charge the nose and alter attitude even when none is individually catastrophic.

3.4.1 A lower-bound fracture-energy screen

A mathematical threshold can be calculated, but it must not be mistaken for a validated breakup diameter. For a linear-elastic brittle surrogate,

$$E' = \frac{E}{1-\nu^2}, G_{Ic} = \frac{K_{Ic}^2}{E'}, E_{section} = G_{Ic} A_p. \quad (9a)$$

Using $K_{Ic} = 5.6$ MPa m❑$^{1/2}$ from the CVD-diamond measurements of Drory et al. [26], $E = 830$ GPa and $v = 0.10$ as an illustrative microcrystalline-diamond surrogate consistent with the property scale reported by Meindlhumer et al. [27], Eq. (9a) gives $G_{Ic} \approx 37.4$ J m^{-2}. For a 50 mm circular section, $E_{section} \approx 0.0734$ J.

If a fraction η_c of a grain's collision energy couples into the specific crack mode, an energy-equivalent grain screen is

$$m_{d,*} = \frac{2 E_{section}}{\eta_c v_r^2}, d_{d,*} = \left(\frac{6 m_{d,*}}{\pi \rho_d} \right)^{1/3}. \tag{9b}$$

For $\rho_d = 2,500$ kg m^{-3}:

Relative speed	Grain at 100% coupling	Grain at 1% coupling
31.6 km s^{-1}	147 ng; 48.2 µm	14.7 µg; 224 µm
100 km s^{-1}	14.7 ng; 22.4 µm	1.47 µg; 104 µm
316 km s^{-1}	1.47 ng; 10.4 µm	147 ng; 48.2 µm
1,000 km s^{-1}	0.147 ng; 4.82 µm	14.7 ng; 22.4 µm

These small thresholds show why dust cannot be ignored; they do **not** show that a grain of the stated size will split the article. $G_{Ic} A_p$ is the minimum surface-creation energy of one ideal through-section crack. It omits impact geometry, crater confinement, attenuation, stress-wave superposition, flaw initiation, crack branching, fragment kinetic energy, dynamic toughening, phase change and the difference between undoped CVD diamond and the proposed BDD architecture. Conversely, an existing flaw or launch-induced tensile field can reduce the local margin. Equation (9b) is therefore a conservative trigger for hydrocode and dynamic-fracture tests, not a ballistic-limit equation.

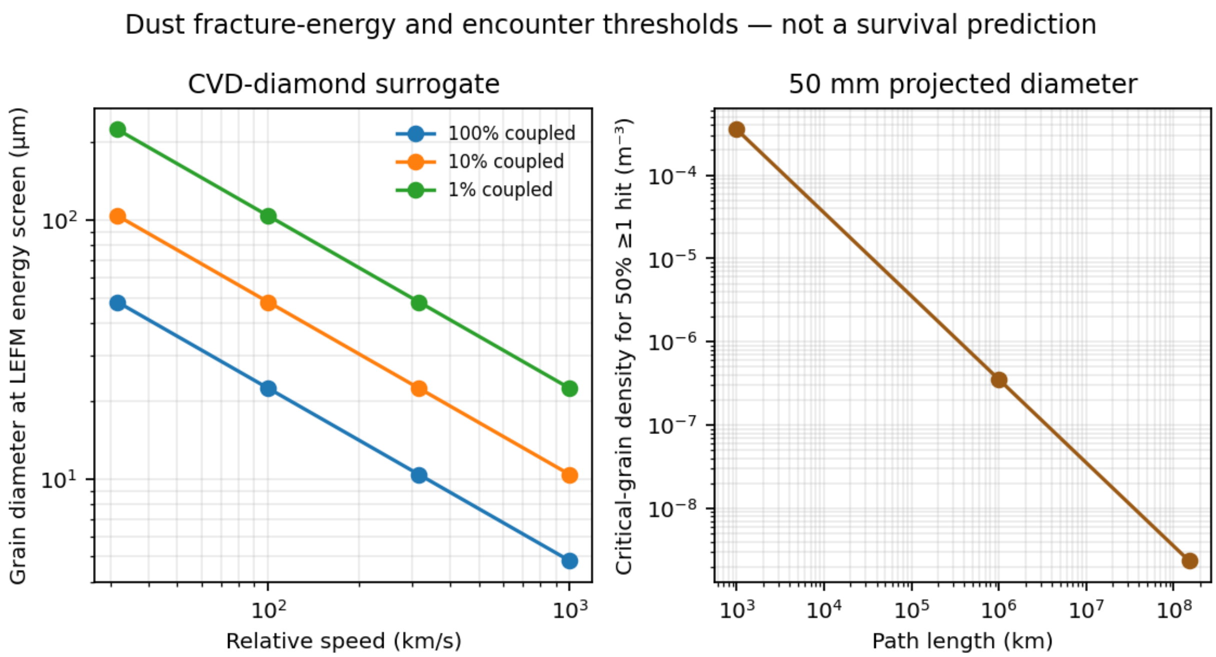

CVD-diamond surrogate fracture-energy and Poisson encounter thresholds. Both panels are parametric screens, not a survival prediction.

3.4.2 Encounter probability over distance

For a directional differential flux integrated along the trajectory, the expected impact count on projected area A_p is

$$\lambda_{\geq m_*} = \int A_p(t) F_{\geq m_*}(r(t), v(t), \Omega, t)\, dt, P_0 = e^{-\lambda_{\geq m_*}}. \quad (10)$$

In the deliberately simplified uniform-number-density case, $\lambda = A_p n_{\geq m_*} D$. The probability of one or more critical hits and the probability of “enough” hits to reach count k are

$$P(N \geq 1) = 1 - e^{-\lambda}, P(N \geq k) = 1 - e^{-\lambda} \sum_{j=0}^{k-1} \frac{\lambda^j}{j!}. \quad (10a)$$

For the 50 mm projected diameter, the cumulative fluence of already-defined critical grains required for 1%, 10%, 50% and 90% probability of at least one hit is respectively 5.12, 53.7, 353 and 1,173 m^{-2}. The equivalent **uniform** critical-grain densities for a 50% chance are:

Path length	$n_{\geq m_*}$ giving 50% chance of at least one hit
1,000 km	3.53×10^{-4} m^{-3}
1,000,000 km	3.53×10^{-7} m^{-3}

Path length	$n_{\geq m_*}$ giving 50% chance of at least one hit
1 au	2.36×10^{-9} m⁻³

This is a sensitivity table, not a claim that space contains those densities. It allows a MEM-class or measured environment output to be tested immediately. At ARGUS speeds the projectile velocity also lies outside ordinary spacecraft use of MEM3, so the underlying directional number-density and velocity distribution must be transformed into the projectile frame rather than blindly reusing a conventional-orbit flux.

Repeated subcritical impacts cannot generally be represented by one universal integer k. A defensible cumulative model assigns a calibrated damage increment ψ_i to each impact:

$$D=\sum_i \psi\left(m_i, v_{r,i}, x_i, \theta_i; M, F\right), P_{fail}=P\left(D \geq 1\right), \quad (10b)$$

where M is the measured material state and F the flaw population. Crater overlap, nose recession, charge, spin impulse and crack coalescence make ψ nonlinear. A publication-quality survival claim therefore requires a trajectory-specific MEM/ORDEM or measured particle distribution, uncertainty bounds, a residual-gas column, impact tests, hydrocode/dynamic-fracture calibration and Monte Carlo propagation. Options to test include a sacrificial conductive or carbon-rich nose, a separated dust shield, in-flight plasma/charge sensing and reduced projected area, but none is credited without showing that it survives launch and preserves the intended impact state.

3.5 Earth atmosphere: energy deposition, not intact arrival

Firing a reference article into Earth's atmosphere is not treated as an intact-impact mode. The sea-level pressure column corresponds to an atmospheric mass per horizontal area of approximately $\Sigma_{atm}=p_0/g \approx 1.03 \times 10^4$ kg m⁻². A 50 mm circular core cross-section would geometrically intercept

$$m_{swept}=A_p \Sigma_{atm} \approx 20.3\,kg, \quad (11)$$

about 68 times the complete 0.300 kg reference moving mass. Retaining a 0.200 m² coupling area would raise the intercepted column above 2,000 kg. The actual flow is not a sequence of perfectly inelastic collisions, but the column comparison alone rules out treating the atmosphere as a small perturbation.

Coupled ablation and radiation models are already required for much slower high-speed sample-return vehicles [22], while meteor research describes sputtering and ablation at collision speeds far below the upper ARGUS range [23]. At 100 km s⁻¹ the reference 0.300 kg article carries 1.5 GJ; at 1,000 km s⁻¹ it carries 150 GJ. Atmospheric molecules striking the nose have about 51.8 eV per atomic mass unit at 100 km s⁻¹ and 5.18 keV per atomic mass unit at 1,000 km s⁻¹. Dissociation, ionisation, radiation, ablation and fragmentation would dominate. The energy would be deposited over an altitude and path distribution as shock-heated gas, plasma, radiation, vapour, fragments and acoustic motion.

Break-up and intense plasma/light production are the expected outcome, but this paper does not assign a single breakup altitude or optical efficiency. Those outcomes require nonequilibrium radiation-hydrodynamics with carbon/boron chemistry, evolving shape, fragmentation and trajectory. The atmosphere is an excluded operational target and a separate scientific energy-deposition problem.

3.6 Reachable-speed evidence ladder

The nearest published macroscopic solid-projectile result located for this audit is an 8 mm, 0.36 g magnesium projectile demonstrated at 10.4 km s^{-1} [24]. A laser-driven experiment has accelerated a microgram-scale plasma projectile to about 140 km s^{-1} [25], but a plasma packet is not evidence that a 0.300-30 kg solid article survives launch or flight.

Speed state	0.300 kg kinetic energy	Evidence interpretation	Dominant unresolved issue
10.4 km s^{-1}	16.2 MJ	solid macroscopic analogue at 0.36 g, not at 0.300 kg	mass scaling and electromagnetic coupling
31.6 km s^{-1}	150 MJ	proposed 0.1% mode; natural-impact analogue only	intact launch and diagnostic containment
100 km s^{-1}	1.5 GJ	T-DART speed or 1% R-DART mode; no solid macroscopic demonstration	ionising dust/gas interaction and fracture
316 km s^{-1}	15 GJ	proposed 10% R-DART mode	high-energy-density impact state
1,000 km s^{-1}	150 GJ	full R-DART accounting point	every solid-survival claim unvalidated

The staged research programme should therefore include a “reachable speed” branch independent of the energy-module branch. A projectile demonstrated at 10.4 km s^{-1} does not validate 31.6 km s^{-1} by interpolation, and a microgram-scale 140 km s^{-1} plasma does not validate a macroscopic solid at 100 km s^{-1}.

4. Scientific and application envelope

4.1 Evidence-labelled utility matrix

Candidate use	Quantity to control	Projectile logic	Evidence status	Principal gate
Hypervelocity shock and equation-of-state research	impact speed and specific energy	R-DART or scaled same-mass mode	established scientific field; ARGUS scale unvalidated	instrumented target, containment and diagnostics
Spacecraft material and shield calibration	speed, geometry and fragment response	purpose-built subscale surrogate	established field; POD-7K extrapolation unvalidated	comparison with existing launchers
Fragmentation and momentum-coupling science	energy, impulse, ejecta and target state	comparative R/V/T campaign	analogue evidence; extreme regime unvalidated	experiments that benchmark hydrocodes

Candidate use	Quantity to control	Projectile logic	Evidence status	Principal gate
Small-body deflection research	impulse, warning and target response	slowest member meeting timing	kinetic impact demonstrated; ARGUS family unvalidated	orbit covariance and fragmentation safety
Airless-body crater or seismic science	controlled local energy and observable response	lowest mode meeting the measurement objective	conceptual for this architecture	science value, planetary protection and site safety
Momentum-beam or pellet-receiver research	repeatable impulse train and receiver coupling	new low-mass projectile	speculative and outside the validated family	receiver survival, pointing and net efficiency
Orbital-debris remediation by impact	post-impact fragment population	no preferred member	benefit not established	demonstrate net debris-risk reduction

The matrix prevents three category errors. Existing impact science is analogue evidence, not evidence for a 150 GJ launcher. A physically calculable outcome is not necessarily a useful or safe application. A projectile family designed around one mass-speed contour need not be optimal for a different scientific instrument.

4.2 Strongest scientific cases

Controlled shock and equation-of-state experiments

A very-high-speed, accurately characterised projectile could create pressure, temperature and phase states that are difficult to reach with conventional macroscopic launchers. The scientific value would come from a measured velocity history, known moving mass, target state, time-resolved diagnostics and recovered or remotely observed response. It would not come from release speed alone.

R-DART is the natural limiting article because it maximises impact speed at the common energy. Lower-energy same-mass modes provide scientifically useful separation of variables. A campaign could compare measured shock observables against hydrocodes and determine where conventional scaling breaks down. Containment may make a full-energy terrestrial target impossible; if so, remote inert space targets and stand-off observation would become part of a separate safety and legal study, not an automatic justification.

Materials, shields and structural response

NASA and ESA facilities demonstrate the value of controlled impacts for spacecraft materials and protection [4-6]. POD-7K is not required for most shield testing, and a 150 GJ macroscopic article would overwhelm rather than calibrate a specimen. The useful contribution is therefore architectural: the force-cell and moving-interface research could produce lower-energy, purpose-built projectiles with controlled geometry and repeatability.

This is also an important falsification route. If a representative multi-cell string cannot outperform simpler launchers in velocity knowledge, repeatability or cost per scientifically useful shot, then the long architecture has no materials-science case.

Momentum coupling, ejecta and fragmentation

The R/V/T comparison provides equal energy with a tenfold momentum range. Instrumented targets of varied porosity, cohesion, volatile content and geometry could test how β, ejecta direction and fragment distribution change with speed. The scientifically important output is a response distribution, not a single best-case coefficient.

This work supports both impact physics and later small-body assessment. It also guards against a common mistake: assuming that the momentum-amplification factor measured by NASA's DART mission can be transferred to a different target or a 100-1,000 km s^{-1} impact [12-14].

4.3 Conditional applications

Small-body intervention. Deflection is an appropriate worked case because it directly uses the momentum ledger and has real kinetic-impact precedent. The outcome is governed by target-specific orbit dynamics and material response, so Section 9 remains a scale screen rather than a mission claim.

Airless-body geophysics. A deliberately low-energy impact on an authorised, scientifically selected airless body might support crater, seismic or subsurface-response studies. The energy should be the lowest that closes the measurement objective. Planetary protection, ejecta evolution, third-party spacecraft risk and scientific opportunity cost would control the decision.

Momentum-beam or pellet-receiver experiments. A sequence of low-mass projectiles could in principle deliver impulse to a purpose-built receiver. That would need a new projectile, receiver, pointing and efficiency model. It is listed as a research question, not as an established use of R-DART or T-DART.

4.4 Uses not established here

High-energy impact is not assumed to remediate orbital debris. It may create a larger and less trackable fragment population, so no family member is recommended without evidence of a net reduction in long-term debris risk.

Terrestrial targets, active spacecraft, coercive demonstrations and combat applications are outside scope. The paper also does not claim cargo launch, because a 100-1,000 km s^{-1} release state is incompatible with treating an ordinary payload as intact merely because the launcher equations close.

5. Functional launcher architecture

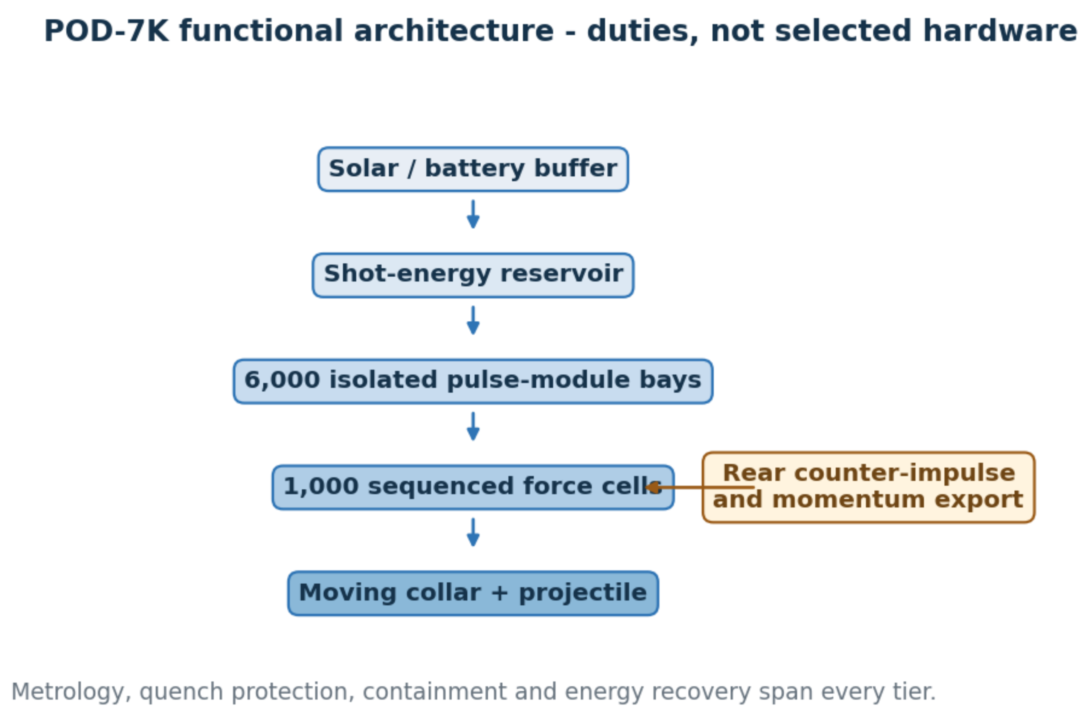

Functional POD-7K architecture. The diagram states energy and momentum duties, not a wiring diagram or selected hardware.

The minimum defensible architecture has five energy-transfer tiers, a moving interface and a separate momentum-export path.

Slow source and battery buffer. Solar generation and a long-duration battery receive energy over hours or days. This tier determines recharge cadence and keeps the pulse inventory electrically isolated during dormancy.

Shot-energy reservoir. A bulk store accepts charge from the buffer and supplies local pulse networks. Its installed energy, deliverable power, fault energy, recovery fraction and thermal state must be measured as one system.

Six thousand module bays. Six fault domains are associated with each of 1,000 force cells. A bay contains isolation, monitoring, a replaceable energy subassembly, switching and a dump path. The bay need not contain one catalogue component.

One thousand force cells. Each seven-metre cell shapes a local waveform against the moving collar. Cell timing, inductance gradient, transverse stability, energy recovery and reaction load require representative tests.

Moving interface. The coupling collar, load spreader, projectile core, guidance and release hardware form one complete moving mass. The 0.300 kg R-DART allocation cannot be transferred to a 30 kg T-DART without a new bill of mass and coupling design.

Rear momentum export. A counter-impulse device is a sixth physical path. It can manage retained centre-of-mass momentum but does not bypass force-cell reaction or structural transmission.

Metrology, quench protection, containment, energy recovery and independent safing span every tier. No tier may be closed by selecting the best public number from unrelated technologies.

6. Six thousand hardware modules and present technology

6.1 A bay count can be physical before its contents close

The earlier papers correctly reject the equation “6,000 energy tranches = 6,000 purchasable batteries”. That does not require 6,000 to remain an accounting fiction. A distributed machine may deliberately provide 6,000 identical module bays while their internal energy assemblies evolve. Sandia’s Z machine uses 36 repeated Marx-bank and intermediate-store channels with multiple capacitors, switches and transmission stages within each channel [8]. Z is not an ARGUS-scale launcher analogue. It is evidence that a pulsed-power “module” can be an isolated chain rather than one capacitor.

For POD-7K, a representative bay must demonstrate:

- a declared share of installed energy and the member-specific power duty;
- bounded fault energy, galvanic isolation and a verified dump path;
- synchronisation with the associated force cell;
- acceptable inductance, voltage reversal and switch life;
- measured transfer efficiency and recoverable energy;
- thermal recovery at the proposed cadence;
- replaceability without disturbing the launcher centreline; and
- a mass and volume extrapolation compatible with deployment.

The defensible statement is this: **6,000 physical module bays are a testable architecture; 6,000 existing 85.714 MJ pulse modules are unsupported.**

6.2 Batteries close slow energy, not firing power

Tesla lists a four-hour Megapack configuration at 3,916 kWh, 979 kW and up to 84,000 lb [9]. On stored energy alone,

$$N_{MP} = \frac{514.286\,GJ}{3,916\,kWh \times 3.6\,MJ\,kWh^{-1}} \approx 36.5\,. \quad (12)$$

Thirty-seven units would exceed the reference nameplate-energy inventory, but their combined published power is about 36.2 MW. The R-DART mean admitted power is 30.612 TW, approximately 8.45×10^{5} times larger. Even 6,000 complete Megapacks would provide only about 5.87 GW at the cited four-hour rating while storing far more than one shot requires.

Such a comparison supports this paper's proposed hierarchy, with an important qualification: Tesla-class batteries can illustrate bulk recharge and state of charge, but a separate pulse reservoir and forming network must supply the shot. The published product is terrestrial and grid-connected; it is not a selected or space-qualified POD-7K component.

6.3 Present capacitors close waveform class but not system mass

General Atomics lists the CMX 3283CMX2640 self-healing capacitor at 255 kJ, about 140 kg and millisecond discharge duty [10]. An equal full-energy module allocation would require

$$N_{CMX/mod}=\frac{85.714\,MJ}{0.255\,MJ}\approx336. \tag{13}$$

That is roughly 47 t of capacitor cans per module before buswork, switches, containment and qualification. Across 6,000 bays, the literal extrapolation is about 2.02 million capacitors and 282,000 t. The same manufacturer lists 100 kV sub-microsecond PDS units with high peak current but only hundreds of joules per unit [11]. They illustrate local edge shaping, not bulk shot energy.

The functionally correct chain is

slow source and battery → bulk pulse reservoir → local forming network → force cell.

It remains unclosed because the complete chain must satisfy energy, power, inductance, voltage, current, isolation, mass, life, heat, recovery and environment together. A future technology may change the answer, but the current paper cannot count that technology before it exists.

7. Electromagnetic and structural definition

The launcher remains a functional reference until a representative cell provides a measured force law. For a switched-inductive or linear-synchronous interpretation,

$$F_x\simeq I_s I_a\frac{dM}{dx}+\frac{1}{2}I_a^2\frac{dL_a}{dx}, \tag{14}$$

where $M(x)$ is mutual inductance and $L_a(x)$ is an effective moving inductance. The 16.410 T Maxwell-pressure value over 0.200 m² is retained only as an ideal force-density screen; it does not replace Eq. (14).

Each repeated force-cell interface must define:

1. geometry, aperture, supports and active length;
2. measured mutual and self-inductance gradients;
3. current, voltage, field and force envelopes;
4. allowable position, timing and transverse errors;
5. collar temperature, strain and retained clearance;
6. admitted energy, kinetic transfer, loss and recovery;
7. normal, late, missing and asymmetric pulse cases;
8. quench, arc and dump behaviour;
9. axial and off-axis reaction delivered to the spine; and
10. inspection and replacement criteria.

The seven-kilometre spine cannot be designed by dividing mean force by a nominal cross-section. A moving force pulse excites modes and changes the metrology state seen by later cells. The development model therefore requires a transient electromechanical thread from one cell to a multi-cell string and then to the complete structure. No unsupported total launcher mass is introduced.

8. Rear counter-impulse and the 1.5 result

Projectile impulse $J_p = p$ gives an uncompensated free-platform velocity change

$$\Delta v_p = \frac{J_p}{M_p}. \tag{15}$$

If a rear device expels mass m_e at effective exhaust speed u_e, its ideal opposed impulse is

$$J_c = m_e u_e. \tag{16}$$

For match fraction $f = J_c / J_p$,

$$\Delta v_{res} = \frac{J_p - J_c}{M_p} = \frac{(1 - f) J_p}{M_p}. \tag{17}$$

These equations validate the valve principle with the required conservation statement. If exhaust is collinear, correctly timed and $f = 1$, the platform's final centre-of-mass velocity change is zero in the ideal model. Momentum has not disappeared: projectile and exhaust carry equal external momentum in opposite directions.

Case	Projectile impulse	Uncompensated platform Δv	Translation during launch	Expelled mass at 4.5 km s^{-1}
R-DART	0.300 MN s	0.150 m s^{-1}	1.05 mm	66.7 kg

Case	Projectile impulse	Uncompensated platform Δv	Translation during launch	Expelled mass at 4.5 km s^{-1}
Q-DART	0.600 MN s	0.300 m s^{-1}	4.20 mm	133.3 kg
V-DART	1.000 MN s	0.500 m s^{-1}	11.7 mm	222.2 kg
M-DART	2.000 MN s	1.000 m s^{-1}	46.7 mm	444.4 kg
T-DART	3.000 MN s	**1.500 m s^{-1}**	**105 mm**	**666.7 kg**

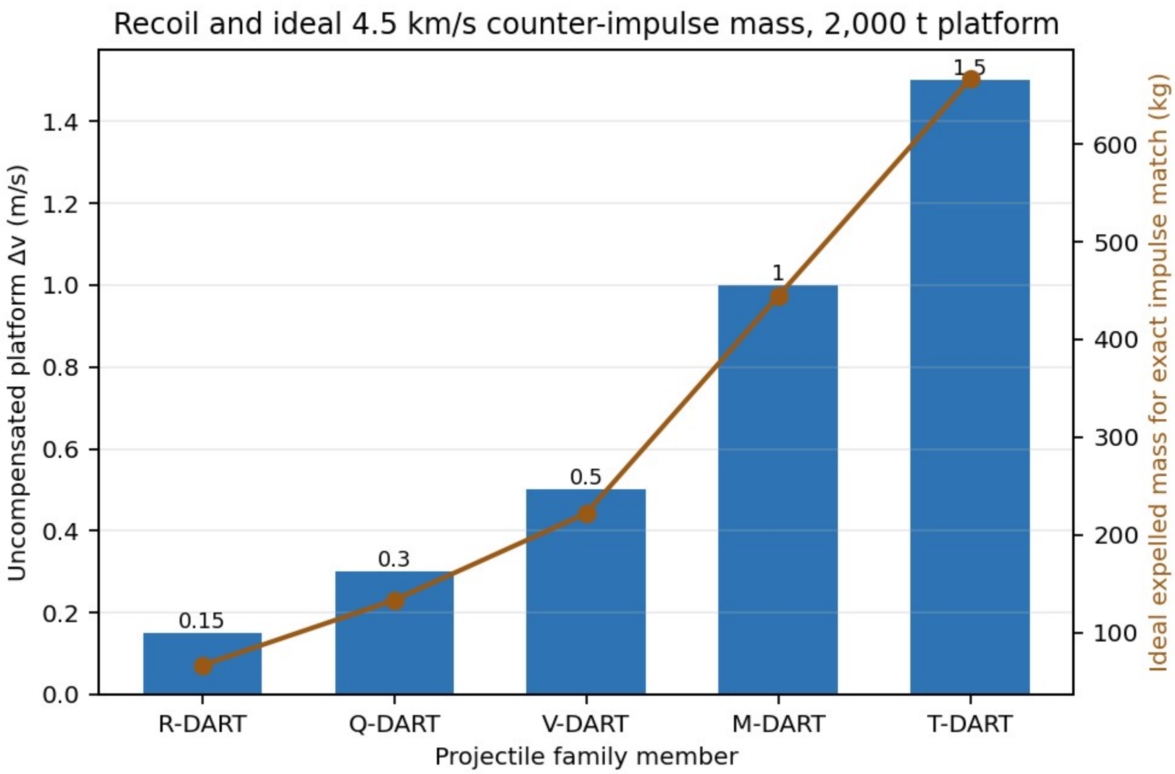

Uncompensated recoil and the ideal 4.5 km s^{-1} exhaust mass required for an exactly opposed impulse.

The proposed 1.5 value is correct for a 3 MN s T-DART shot and a 2,000 t platform, with units of metres per second. It is not the total displacement. During a 140 ms constant-force launch, the ideal uncompensated translation is about 0.105 m; after launch the 1.5 m s^{-1} drift continues until momentum is restored.

If the allowable residual speed is v_{tol}, the ideal fractional impulse error must satisfy

$$|1-f| \le \frac{M_p v_{tol}}{J_p}. \tag{18}$$

For T-DART, a 2,000 t platform and a 0.1 mm s^{-1} tolerance, the match must be within about 6.7×10^{-5} before direction and torque error are added. A fast valve using a pressurised reservoir might supply a short impulse but must declare exhaust speed, mass, plume and throttling. A high-specific-impulse electric thruster might use less mass but cannot be assumed to produce tens of meganewtons over milliseconds; it may instead restore momentum slowly while the platform drifts.

The counter-impulse does not remove force-cell reaction, axial load between launcher and rear device, bending from line-of-action error, plume hazards, expellant logistics or post-shot metrology. Its engineering value is the control of retained platform translation, not the erasure of recoil physics.

9. Worked application: generic small-body impulse screen

9.1 Assumptions and selection rule

NASA describes indicative 50 m, 140 m and 1,000 m asteroid hazard classes [12]. This deliberately simplified example uses spherical targets at 50, 140, 300 and 1,000 m, rocky density $\rho_a = 2{,}600$ kg m^{-3}, inbound speed $u_a = 20$ km s^{-1} and an illustrative $\beta = 3$. DART constrained momentum amplification for Dimorphos, but its 2.2-4.9 result is not a validated constant for other materials or the present impact-speed range [13-15].

$$M_a = \frac{\pi}{6} \rho_a D_a^3, \Delta v_a = \frac{\beta p}{M_a}. \quad (19)$$

For a generic intercept distance R, the minimum one-dimensional warning and ideal transverse shift are

$$T_{warn,min} = \frac{R}{u_a} + \frac{R}{v}, s_{lin} = \Delta v_a \frac{R}{u_a}. \quad (20)$$

The screen omits solar gravity, the true orbit, Earth focusing, b-plane geometry, rotation, nonlinear repeated impacts and fragmentation. NASA/JPL's NEO Deflection App illustrates why a real decision must use orbit-specific geometry [16].

The family selection rule is:

> *Reject every projectile that cannot reach the declared intercept distance within the warning window. Among the admissible articles, choose the slowest for maximum direct impulse unless target-response or safety evidence favours a different speed.*

This is circumstance-based projectile selection. It applies to a declared objective and geometry, not to "asteroids" as a uniform target category.

9.2 Scale results

The following table uses T-DART, $\beta = 3$, the declared density and a screening displacement of two mean Earth radii.

Target diameter	Intercept distance	Minimum warning	Single-shot linear shift	Ideal T-DART shots for two Earth radii
50 m	0.1 AU	10.4 d	39.6 km	323
50 m	1 AU	103.9 d	395.6 km	33
50 m	5 AU	519.4 d	1,978 km	**7**
140 m	0.1 AU	10.4 d	1.80 km	7,071
140 m	1 AU	103.9 d	18.0 km	708
140 m	5 AU	519.4 d	90.1 km	**142**
300 m	0.1 AU	10.4 d	0.183 km	69,573
300 m	1 AU	103.9 d	1.83 km	6,958
300 m	5 AU	519.4 d	9.16 km	**1,392**
1,000 m	0.1 AU	10.4 d	0.00494 km	2,576,747
1,000 m	1 AU	103.9 d	0.0494 km	257,675
1,000 m	5 AU	519.4 d	0.247 km	**51,535**

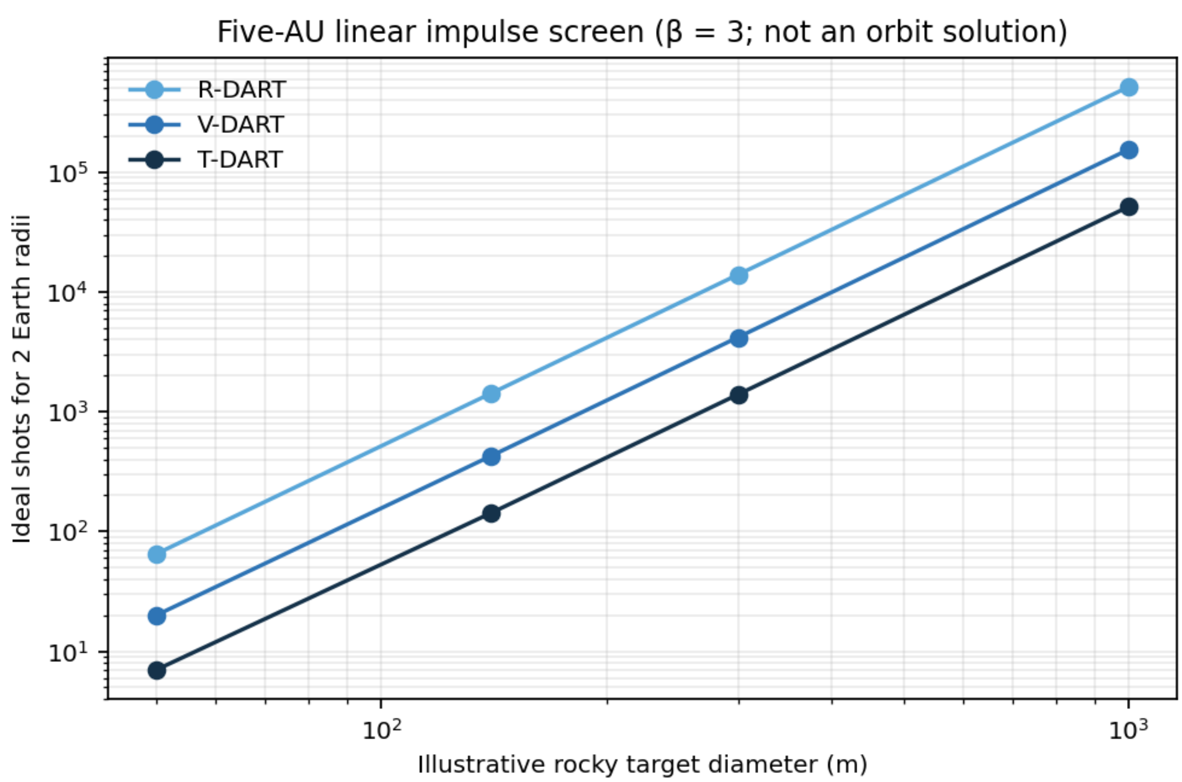

Ideal shot count versus target diameter at 5 AU. The logarithmic figure shows scale and is not a mission-effectiveness prediction.

The 50 m, 5 AU case is the only listed screen in a single-digit regime. It would still require validated repeated impacts, target-specific navigation, cadence, plume-safe momentum management and evidence that fragmentation does not make the outcome worse. A 140 m body requires 142 ideal shots even at 5 AU; 300 m and 1 km bodies require thousands to tens of thousands. The current 150 GJ family is therefore not a defensible single-shot solution to a large asteroid.

This negative result strengthens the broader paper. It demonstrates why the launcher should be assessed first as an experimental platform and why projectile effectiveness must be stated for a circumstance, not advertised from speed alone.

10. Development and falsification programme

Gate	Article or model	Required result	Stop or redesign condition
S0	existing-facility benchmark	define the experiment that POD-7K uniquely improves	no distinct scientific value
L1	module-bay emulator	isolation, switching, dump and timing at scaled energy	uncontrolled fault propagation
L2	representative pulse module	joint energy-power-life performance	mass or heat incompatible with deployment
L3	force cell	measured force law, efficiency, stability and reaction	no scalable coupling topology
L4	multi-cell string	phasing, recovery and transient structural response	clearance or waveform cannot be preserved
P1	weighed moving article	complete mass, release, stability and retained geometry	mass or release ledger fails
X1	instrumented impact	velocity, target state and diagnostic closure	response cannot be reconstructed
X2	comparative R/V/T campaign	separate speed, momentum and material effects	no transferable response model
P2	matched material coupons	diamond, BDD, diamond-SiC, SiC, B4C, W and carbon-rich comparisons	no architecture retains its declared state
E1	dust-impact and residual-gas article	material-specific crater, plasma, charge and dynamic-fracture thresholds	no credible flight-survival envelope
E2	trajectory environment model	particle-frame fluence, Poisson critical-hit and compound cumulative-damage probability	probability exceeds the declared mission limit
E3	atmospheric energy-deposition model	breakup, ablation, plasma and radiation distribution	terrestrial passage retained as a claimed intact mode
R1	rear counter-impulse rig	impulse magnitude, vector, timing and plume model	residual or torque exceeds budget

Gate	Article or model	Required result	Stop or redesign condition
R2	cell plus counter-impulse	net translation suppressed while reaction is measured	structural load remains uncontrolled
M1	orbit and covariance model	target-specific b-plane effect	linear screen overstates effect
M2	impact campaign and hydrocodes	response distribution across target classes	fragmentation uncertainty dominates
D1	unarmed space demonstrator	charging, safing, metrology and momentum reconstruction	no lawful safe operating envelope

The sequence starts with scientific value rather than scale. A project should not build a larger pulse system until it can name the measurement that the next stage makes possible. Energy progression, moving-article progression and evidence progression must remain separate controls.

The following results would falsify the central proposition:

- no representative force-cell topology achieves measurable, repeatable coupling at an extrapolatable efficiency;
- synchronisation error or moving-interface instability grows faster than the multi-cell architecture can correct;
- deployable energy-storage mass cannot satisfy the combined energy, power, life and thermal duty;
- target diagnostics cannot reconstruct the impact state well enough to distinguish the projectile-family hypotheses;
- counter-impulse error or plume risk exceeds the platform-control benefit;
- no candidate or graded material architecture survives the measured acceleration-plus-dust state with adequate statistical confidence;
- particle fluence makes critical-hit or accumulated-damage probability unacceptable over the required path; or
- an existing, smaller launcher provides the same scientific information with substantially lower risk and cost.

11. Conclusions

The launcher and projectile system has a broader and more defensible research case than a paper organised solely around asteroid interception.

1. The 150 GJ family is an **experimental control surface**. R-DART maximises speed and access to extreme shock states; T-DART maximises direct impulse; V-DART provides an important intermediate comparison.

2. A variable-energy sequence is essential. For the 0.300 kg article, 0.1%-100% mechanical energy spans 31.6-1,000 km s^{-1} and reduces mean module power from 5.102 GW to 0.161 MW at the lowest listed mode. These are accounting points, not proven operating modes.
3. The strongest independent scientific utilities are controlled hypervelocity impact, shock and equation-of-state research, comparative momentum coupling, fragmentation science and purpose-built subscale spacecraft-material testing.
4. Space is not a complete vacuum. Dust and residual gas become energetic impactors in the projectile frame. At 1,000 km s^{-1} a 1 µg grain supplies 500 J locally. The CVD-diamond LEFM result is only a lower-bound crack-surface-energy screen, not a ballistic limit; the Poisson calculation shows how a trajectory-specific critical-grain density or fluence becomes a hit probability.
5. No material is selected. Boron-doped diamond is not a single property set, and high boron concentration can reduce elastic constants, hardness and thermal conductivity. CVD diamond, diamond-SiC, toughened SiC, boron carbide, tungsten and carbon-rich sacrificial layers remain comparative candidates. The full R-DART acceleration gives a 5.49 GPa one-dimensional end-loaded diamond-core stress screen before dynamic and flaw effects.
6. An Earth-atmosphere trajectory is an energy-deposition event, not an intact impact. A 50 mm cross-section encounters about 20.3 kg of atmospheric column, so ablation, ionisation, radiation, plasma and fragmentation must dominate the analysis.
7. Six thousand can be a physical module-bay count. It cannot yet be a claim that 6,000 complete 85.714 MJ pulse units exist. Tesla-class batteries fit a slow buffer tier; present catalogue capacitors do not close full-system mass.
8. The rear valve can suppress retained platform translation only by exporting equal momentum. For T-DART and a 2,000 t platform, 1.5 m s^{-1} is the correct uncompensated recoil speed and 666.7 kg at 4.5 km s^{-1} is the ideal counter-impulse mass.
9. Asteroid deflection remains a useful worked example but not a universal justification. The present linear screen rejects a single-shot large-asteroid claim and shows why target, distance, warning and desired outcome must be stated.
10. Airless-body geophysics and momentum-beam research are conditional future cases. Orbital-debris impact, terrestrial targets, active satellites and coercive use are not supported here.

The proper next step is not a full-scale launcher. It is a scientific qualification programme that proves one pulse bay, one force cell, one weighed moving article and one instrumented target response, then tests whether those results scale. That programme can generate valuable impact science even if the complete POD-7K architecture is ultimately falsified.

Declarations

Funding. No external funding is declared.

Competing interests. The author declares no competing financial interest.

Data and code availability. The deposited reproducibility package contains the reference-family, operating-mode, module-duty, counter-impulse, application-utility, material-evidence, dust-collision, dust-fracture, dust-encounter, flight-environment and generic small-body screens, unit tests and figures used here. It contains no target ephemerides, firing logic, trajectory optimiser, electromagnetic finite-element files or buildable pulse-power design.

Dual-use boundary. The analysis is limited to civil scientific research, authorised inert targets, planetary science and generic planetary-defence screening. Terrestrial, active-satellite, coercive and combat uses are excluded.

References

1. Farmer, C. I. V. *ARGUS-PD/POD-7K Paper I: Parametric Design-Space Analysis of a Modular Projectile Family and Electromagnetic Launch Interface*. Version 1.1.0, companion preprint (2026).
2. Farmer, C. I. V. *ARGUS-PD/POD-7K Paper II: Integrated Architecture and Feasibility Bounds for a Fixed Seven-Kilometre Electromagnetic Planetary-Defence Research System*. Version 1.1.0, companion preprint (2026).
3. Farmer, C. I. V. *ARGUS-PD/POD-7K Paper III: Coupled Feasibility Boundaries, Uncertainty Propagation and Falsification Criteria*. Version 1.1.0, companion preprint (2026).
4. NASA Johnson Space Center. Hypervelocity Impact Technology: facility, impact-physics and materials-testing overview (accessed 29 July 2026). https://hvit.jsc.nasa.gov/.
5. Christiansen, E. L. et al. *Impact Testing for Materials Science at NASA*. NASA Technical Reports Server, 2011. https://ntrs.nasa.gov/api/citations/20110010224/downloads/20110010224.pdf.
6. European Space Agency. Hypervelocity impacts and protecting spacecraft: laboratory testing and shield assessment (accessed 29 July 2026). https://www.esa.int/Space_Safety/Space_Debris/Hypervelocity_impacts_and_protecting_spacecraft.
7. Asay, J. R. & Hawke, R. S. *A Survey of Hypervelocity Impact Technology and Applications*. U.S. Department of Energy report. https://www.osti.gov/servlets/purl/5370387.
8. Sandia National Laboratories. How the Z Machine Works: Marx banks, intermediate stores and triggered switches (accessed 29 July 2026). https://www.sandia.gov/z-machine/about-z/how-z-works/.
9. Tesla. Megapack product details: 3,916 kWh, 979 kW four-hour configuration and 84,000 lb maximum weight (accessed 29 July 2026). https://www.tesla.com/megapack/design.
10. General Atomics. Series CMX self-healing energy-storage capacitors: 3283CMX2640, 255 kJ and approximately 140 kg (accessed 29 July 2026). https://www.ga.com/capacitors/series-cmx-self-healing-energy-storage-capacitors.
11. General Atomics. Series PDS/PDSS fast-pulse capacitors: high-voltage, sub-microsecond pulse ratings (accessed 29 July 2026). https://www.ga.com/capacitors/series-pds-pdss-fast-pulse-capacitors.
12. NASA Science. Asteroid facts: indicative hazard classes at 50 m, 140 m and 1,000 m (accessed 29 July 2026). https://science.nasa.gov/solar-system/asteroids/facts/.
13. NASA. DART data validates kinetic impact as a planetary-defence method: momentum amplification and warning-time discussion (1 March 2023). https://www.nasa.gov/science-research/planetary-science/nasas-dart-data-validates-kinetic-impact-as-planetary-defense-method/.
14. Cheng, A. F. et al. Momentum transfer from the DART mission kinetic impact on asteroid Dimorphos. *Nature* 616, 457-460 (2023). https://doi.org/10.1038/s41586-023-05878-z.

15. Raducan, S. D. et al. Influence of projectile geometry on momentum transfer and implications for DART. *International Journal of Impact Engineering* 162, 104147 (2022). https://doi.org/10.1016/j.ijimpeng.2021.104147.
16. NASA/JPL Center for Near-Earth Object Studies. NEO Deflection App: impulsive asteroid-deflection mission screening (accessed 29 July 2026). https://cneos.jpl.nasa.gov/nda/.
17. Moorhead, A. V., Kingery, A. & Ehlert, S. NASA's Meteoroid Engineering Model 3 and its ability to replicate spacecraft impact rates. *Journal of Spacecraft and Rockets* 57 (2020). NASA report and model documentation: https://ntrs.nasa.gov/citations/20200000563.
18. Lee, N. & Close, S. Meteoroid impact detection for exploration of asteroids: small satellites for asteroid characterization. *Journal of Spacecraft and Rockets* 55 (2018). https://doi.org/10.2514/1.A33928.
19. Christoffersen, R. et al. Experimental hypervelocity dust impacts in olivine: microcraters, melt, vapour and plasma-plume evidence. *Journal of Geophysical Research: Planets* (2026). https://ntrs.nasa.gov/api/citations/20260002134/downloads/ChristoffersenEtAl-Microcraters_JGR_FinalPublished.pdf.
20. Liu, X. et al. Elastic and mechanical softening in boron-doped diamond. *Scientific Reports* 7, 42921 (2017). https://doi.org/10.1038/srep42921.
21. Prikhodko, D. et al. Thermal conductivity of synthetic boron-doped single-crystal HPHT diamond from 20 to 400 K. *MRS Communications* 6, 71-76 (2016). https://doi.org/10.1557/mrc.2016.12.
22. Johnston, C. O. Study of aerothermodynamic modeling issues relevant to high-speed sample return vehicles. NATO STO-AVT-218 (2014). https://ntrs.nasa.gov/citations/20140004052.
23. Guttormsen, G., Fletcher, A. C. & Oppenheim, M. M. Atomic-scale simulations of meteor ablation. *Journal of Geophysical Research: Space Physics* 125, e2020JA028229 (2020). https://doi.org/10.1029/2020JA028229.
24. Huneault, J. et al. An explosively driven launcher capable of 10 km s^{-1} projectile velocities. *Shock Waves* 32, 569-591 (2022). https://doi.org/10.1007/s00193-022-01095-1.
25. Badziak, J. et al. Efficient acceleration of a dense plasma projectile to hyper velocities in the laser-induced cavity pressure acceleration scheme. *Scientific Reports* 7, 9491 (2017). https://doi.org/10.1038/s41598-017-10075-0.
26. Drory, M. D., Dauskardt, R. H., Kant, A. & Ritchie, R. O. Fracture of synthetic diamond. *Journal of Applied Physics* 78, 3083-3088 (1995). https://doi.org/10.1063/1.360060.
27. Meindlhumer, M. et al. Micromechanical properties of micro- and nanocrystalline CVD diamond thin films with gradient microstructures and stresses. *Journal of Vacuum Science & Technology A* 42, 023401 (2024). https://doi.org/10.1116/6.0003235.
28. Ast, J. et al. Micro-cantilever testing of diamond-silicon carbide interfaces in silicon carbide bonded diamond materials produced by reactive silicon infiltration. *Open Ceramics* 8, 100176 (2021). https://doi.org/10.1016/j.oceram.2021.100176.

29. Chen, D., Gilbert, C. J., Zhang, X. F. & Ritchie, R. O. High-temperature cyclic fatigue-crack growth behavior in an in situ toughened silicon carbide. *Acta Materialia* 48, 659-674 (2000). https://doi.org/10.1016/S1359-6454(99)00406-1.
30. Zhang, H. et al. Impact resistance of boron carbide ceramics from hypersonic and supersonic impacts: a large-scaled molecular dynamics simulation study. *Journal of Applied Physics* 131, 135105 (2022). https://doi.org/10.1063/5.0080252.
31. Gludovatz, B., Faleschini, M., Wurster, S., Hoffmann, A. & Pippan, R. Influence of microstructure on the fracture toughness of tungsten alloys. *TMS Annual Meeting Proceedings*, 449-454 (2008). https://pure.unileoben.ac.at/ws/files/1008984/Influence_of_Microstructure_on_the_Fracture_Toughness_of_Tungsten_Alloys.pdf.
32. Iyer, K. A., Mehoke, D. S. & Batra, R. C. Interplanetary dust particle shielding capability of spacecraft multilayer insulation. *Journal of Spacecraft and Rockets* 52, 584-594 (2015). https://doi.org/10.2514/1.A33101.
33. Hoang, T., Lazarian, A., Burkhart, B. & Loeb, A. The interaction of relativistic spacecrafts with the interstellar medium. *Astrophysical Journal* 837, 5 (2017). https://doi.org/10.3847/1538-4357/aa5da6.