DEFENSE & SYSTEMS ECONOMICS

The Economics of Multi-Domain Warfare

Why Coordination Architecture, Not Platform Superiority, Determines the Cost of Modern Combat

Sinéad O'Sullivan

IDEA IN BRIEF

THE ASSUMPTION Advanced militaries pursue multi-domain operations—the simultaneous coordination of air, land, sea, space, cyber, and electromagnetic capabilities—on the premise that integration across domains produces decisive advantage. A force that can see, strike, and maneuver across all domains simultaneously should overwhelm one that cannot.

THE PROBLEM Multi-domain integration generates coordination costs that scale non-linearly with the number of domains, actors, and interfaces involved. Each additional domain does not simply add capability; it multiplies the coordination burden on the architecture that must synchronize them. Against asymmetric adversaries who exploit the seams between domains rather than competing within them, the cost structure inverts: the integrating force bears the architecture cost, while the disruptive force targets the cheapest interface.

THE IMPLICATION The economics of multi-domain warfare are not determined by platform superiority or aggregate spending. They are determined by the ratio between coordination cost and disruption cost. When coordination is expensive and disruption is cheap, the architecture becomes the vulnerability, not the asset.

The doctrine of multi-domain operations has become the organizing principle of advanced military strategy. The United States, NATO, China, and an expanding list of middle powers now structure their force development, procurement, and training around the premise that decisive advantage flows from the ability to coordinate effects simultaneously across air, land, sea, space, cyber, and the electromagnetic spectrum. The logic is that a force which can integrate capabilities across all these domains should see more, decide faster, and strike harder than one confined to fewer dimensions.

What this doctrine does not adequately address is the cost structure it creates: Multi-domain integration is not free, as it generates coordination costs—in data architecture, identification systems, doctrinal harmonization, standards interoperability, command structures, and in the institutional machinery required to make heterogeneous systems behave as a coherent whole. These costs scale non-linearly. Each additional domain, each additional coalition partner, each additional platform type does not simply add capability, but multiplies the number of interfaces that must be managed, the number of potential failure modes that must be anticipated, and the coordination burden on the architecture that holds the system together.

Against symmetric adversaries operating a comparable multi-domain architecture, these costs are mutual in that both sides pay them. But against asymmetric adversaries—those who cannot or choose not to match the integrating force domain for domain—the economics shift dramatically.

In short, the asymmetric adversary does not need to build a coordination architecture; it just needs to disrupt one.

This creates an obvious cost inversion that has profound implications for defense spending, force design, and strategic competition. The multi-domain integrator pays the full cost of synchronization across every interface whereas the asymmetric disruptor pays only the cost of exploiting the weakest interface. The economics of this exchange are structurally unfavorable to the integrator—and they become more unfavorable as the number of domains increases.

THE COORDINATION COST OF INTEGRATION

To understand why, it helps to think about multi-domain operations as an architecture problem rather than a platform problem.

In a single-domain force—an air force operating independently, a navy conducting maritime operations in isolation—the coordination architecture is relatively simple. Platforms share common data links, operate under a unified command structure, follow shared doctrine, and are designed to the same technical standards. Thus, the interfaces between components are internal to the domain and can be optimized through decades of institutional development.

Multi-domain operations break this simplicity. When an air force must coordinate with ground-based air defense, naval strike groups, space-based intelligence, cyber operations, and electronic warfare simultaneously, the number of interfaces between systems grows combinatorially. With two domains, there is one interface. With three, there are three. With six domains, there are fifteen pairwise interfaces—each requiring its own data protocols, timing synchronization, identification procedures, and rules of engagement.

But the real cost is not in the pairwise interfaces. It is in the *multi-way interactions*: situations where an action in one domain affects the operating conditions in two or three others simultaneously. An electronic warfare operation that jams enemy radar also degrades friendly identification systems. A cyber operation that disrupts adversary command networks may affect the timing of a coordinated air-ground strike. A space-based sensor that provides targeting data to naval forces must be protected by air defense assets whose engagement zones may conflict with friendly aircraft transit corridors.

These interaction effects do not scale linearly with domain count, they scale *combinatorially*. And each interaction requires coordination architecture—doctrine, procedures, data links, standards, training, and institutional capacity—to manage.

The coordination architecture is, in economic terms, the fixed infrastructure cost of multi-domain operations. It must be built, maintained, exercised, and updated regardless of whether any particular operation requires all domains simultaneously.

STIGLER IN THE BATTLESPACE

The economist George Stigler observed in 1951 that the division of labor is limited by the extent of the market, and that specialization is only viable when demand is large and stable enough to sustain dedicated producers. The same principle applies to military coordination architecture, but in reverse.

In multi-domain operations, the “market” is the operational environment: the set of scenarios, threat types, and coalition configurations that the architecture must serve. The more diverse this environment—the more domains, the more partners, the more threat types—the more the coordination architecture must generalize. However, generalization is expensive. An architecture optimized for air-land integration may perform poorly in air-maritime-cyber operations, just as an architecture designed for bilateral U.S.-allied operations may fragment when extended to a five-nation coalition with heterogeneous equipment.

The result is a structural tension. The doctrine demands maximum integration across all domains, whereas the economics demand specialization. In other words, the architecture that attempts to do everything across all domains for all partners simultaneously is the architecture that does nothing cheaply.

This is the defense equivalent of the architecture-lag problem that prevents frontier technologies from becoming industries.

The capability at each node—the fighter, the satellite, the cyber tool, the ship—may be excellent, but the system that must coordinate them bears costs that grow faster than the capability it enables.

Importantly, at some point, the marginal domain added to the architecture costs more to coordinate than the capability it contributes. The force therefore becomes more expensive to operate without becoming proportionally more effective.

THE ASYMMETRIC COST INVERSION

Now consider the adversary who faces this architecture. The conventional framing treats asymmetric warfare as a problem of capability mismatch: the weaker side lacks the platforms, training, or technology to compete head-to-head. However this framing is incomplete, as in a multi-domain environment, the asymmetric adversary has a

structural economic advantage that is independent of its platform quality.

The integrating force must synchronize every interface. But the disruptor must only find one that fails.

The cost structure of this exchange is radically asymmetric. Consider a coalition operating across six domains with fifteen pairwise interfaces. The coalition must invest in coordination architecture for every interface: data links, identification-friend-or-foe (IFF), doctrine, standards, timing, and decision procedures. Each interface has a cost of establishment, a cost of maintenance, and a cost of degradation management—the investment required to keep the interface functioning when the adversary attempts to disrupt it.

The adversary's targeting problem is simpler. It does not need to defeat the coalition across all domains, it only needs to identify the interface with the lowest disruption cost—the seam in the architecture where a relatively cheap action produces disproportionate system-level effects—and exploit it.

A GPS jammer that degrades timing synchronization across multiple domains; an electronic warfare system that creates identification ambiguity in one engagement zone; or a swarm of cheap drones that saturates one air defense sector, forcing the architecture to reallocate coordination bandwidth from other interfaces.

Each of these is cheap relative to the architecture it targets, and each attacks not a platform but an *interface*—the connection between platforms that the coordination architecture must maintain.

This is the cost inversion.

The more domains the integrating force adds, the more interfaces it must defend. The more interfaces it defends, the larger its coordination architecture. And the larger the architecture, the more targets the disruptor has, with a greater the probability that at least one interface can be exploited cheaply.

The integrator's cost curve rises with domain count, but the disruptor's cost curve stays flat or falls, because it needs only to find the cheapest entry point in an expanding attack surface.

INTERFACE	INTEGRATOR'S COST	DISRUPTOR'S COST
Data links	Must connect every platform across every domain in real time with compatible protocols	Needs to jam, spoof, or degrade one link to fragment the picture
IFF / identification	Must positively identify every object in a congested multi-domain battlespace	Needs to create ambiguity in one engagement zone to induce hesitation or fratricide
Doctrine	Must harmonize rules of engagement, engagement authority, and escalation procedures across coalition partners	Operates under a single command with simple, permissive ROE
Standards	Must maintain interoperable technical standards across platforms, nations, and domains	Needs only to exploit the gap where two standards fail to interoperate
Timing	Must synchronize actions across domains to a shared clock under degraded conditions	Needs to disrupt synchronization in one domain to desynchronize the whole
Decision speed	Must process multi-domain information faster than the threat evolves	Needs only to present threats faster than the architecture can classify them

Table 1. The cost asymmetry of multi-domain coordination. At every interface, the integrating force bears the full cost of synchronization. The disruptor bears only the cost of exploiting the weakest link.

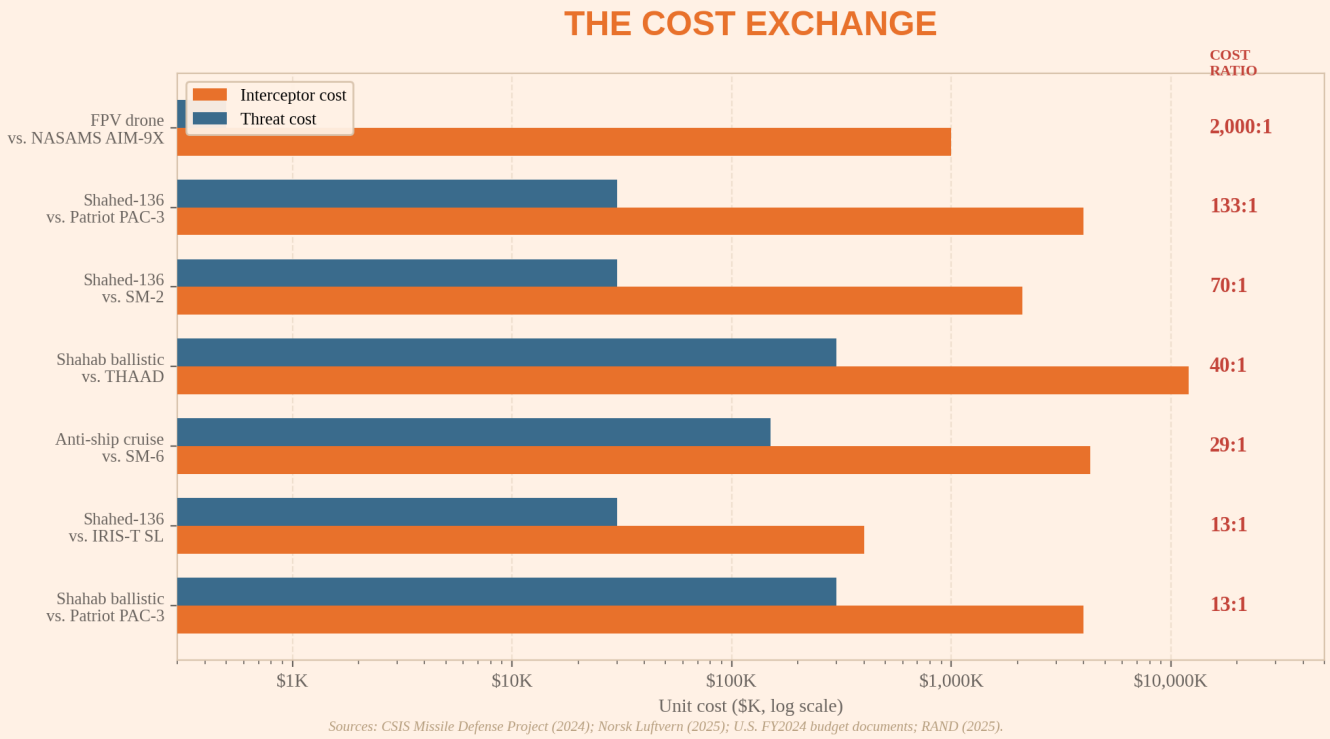

Figure 1. *Unit cost of threat vs. unit cost of interceptor across real engagements. In every pairing, the defender pays more—often by orders of magnitude. When multiple interceptors are required per threat, the ratio worsens further. Sources: CSIS Missile Defense Project (2024); Norsk Luftvern (2025); U.S. FY2024 budget documents; RAND (2025).*

THE FRIENDLY FIRE TAX

The cost inversion has a second, less discussed dimension: the cost the integrating force imposes on itself. When coordination architecture is stretched beyond its capacity, the system does not merely become less effective, it actually becomes dangerous to its own participants.

Friendly fire—fratricide—is not a random occurrence in multi-domain operations; it is a structural consequence and predicted failure mode of coordination overload. When the architecture governing identification, classification, and engagement authority cannot process the multi-domain battlespace at the speed the environment demands, the system produces classification errors, for example an aircraft returning from a strike mission is misidentified as a threat; or a ground unit in a contested zone is engaged by friendly air support operating on a different data picture

These events have occurred in every major coalition air campaign since 1991. And they are not declining in frequency with better technology, because the technology is not the binding constraint. The binding constraint, instead, is the coordination architecture, and the complexity of the environment it must manage is growing faster than the architecture's capacity to process it.

From an economic perspective, fratricide is a coordination cost borne entirely by the integrating force. It destroys friendly capital—aircraft, vehicles, personnel—while the adversary expends nothing. It degrades trust between coalition partners, increasing the institutional friction that slows future coordination. And it generates political costs that constrain operational freedom, reducing the effective scope of the architecture.

The adversary does not need to cause fratricide deliberately, although some asymmetric strategies are explicitly designed to create the conditions for it.

Electronic warfare that degrades identification systems, drone swarms that saturate engagement zones, decoys that crowd the air picture—all of these increase the probability that the coordination architecture will fail and the integrating force will shoot itself.

Thus, the asymmetric adversary's most cost-effective weapon may not be one that destroys enemy platforms directly; it may be one that causes the enemy's architecture to destroy its own.

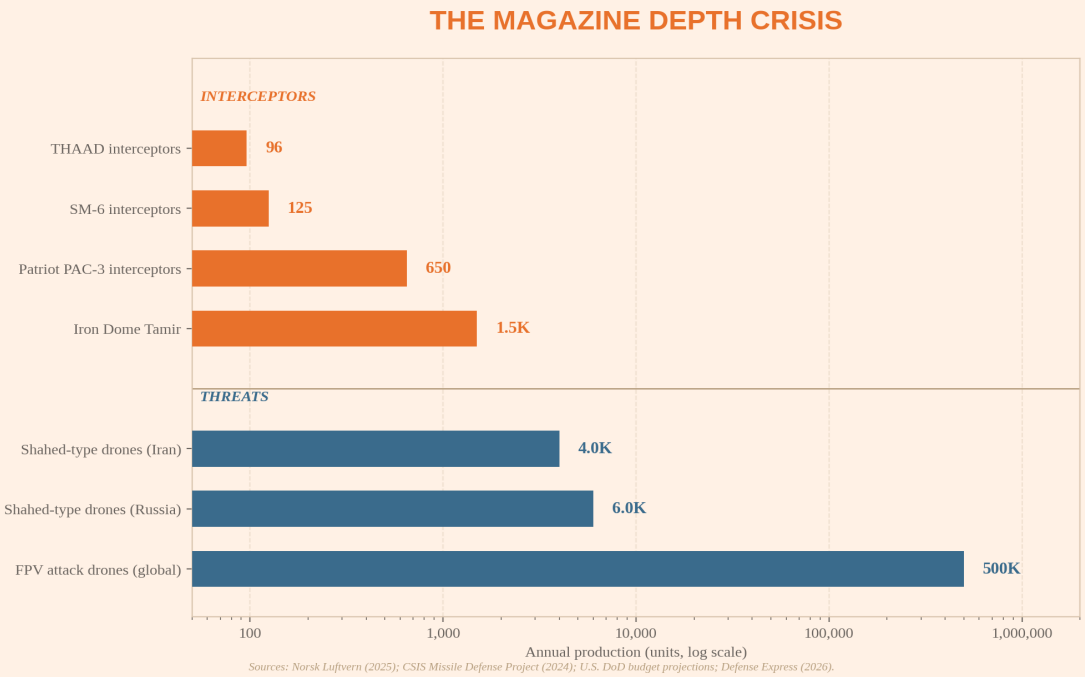

Figure 2. *Annual global production capacity for key interceptor systems vs. the attack drones and missiles they must counter. The production asymmetry compounds the cost asymmetry: defenders cannot manufacture interceptors fast enough to match the volume of threats. Scale is logarithmic. Sources: Norsk Luftvern (2025); CSIS Missile Defense Project (2024); U.S. DoD budget projections; Defense Express (2026).*

THE ARCHITECTURE BANDWIDTH CONSTRAINT

The framework of adaptive bandwidth—developed in O'Sullivan's prior work on institutional coordination—provides a useful lens here. Adaptive bandwidth measures the rate at which an institutional architecture can process, classify, and act on new information.

In economic development, low-bandwidth institutions cannot keep pace with the demands of rapid industrialization. In military operations, low-bandwidth coordination architectures cannot keep pace with the speed and complexity of multi-domain combat.

The relevant bandwidth in multi-domain operations is not data throughput, as modern networks can transmit enormous volumes. It is rather a *decision* bandwidth: the rate at which the system can convert raw sensor data across multiple domains into coherent, correct, and timely decisions. This requires not only data transmission but data fusion, cross-domain correlation, identification resolution, doctrinal application, and command authorization. Each step involves either human judgment or automated logic that must be consistent across every national system participating in the architecture.

When decision bandwidth is exceeded—when the coordination load surpasses the architecture's capacity to process it—the system degrades, meaning that classification errors increase, engagement timelines compress beyond the point where cross-checks are possible, and subsystems begin making independent decisions based on incomplete pictures. The architecture hence fragments from a coherent system into a collection of nodes acting on partial information.

The asymmetric adversary's optimal strategy is therefore not to destroy the architecture's physical components, but to exceed its bandwidth and to present more simultaneous stimuli than the architecture can classify, creating ambiguity faster than the system can resolve it.

This will ultimately force the coordination burden above the threshold where the architecture produces reliable outputs. The result therefore is not a single point of failure but an entire systemic degradation in which the architecture's own complexity becomes its vulnerability.

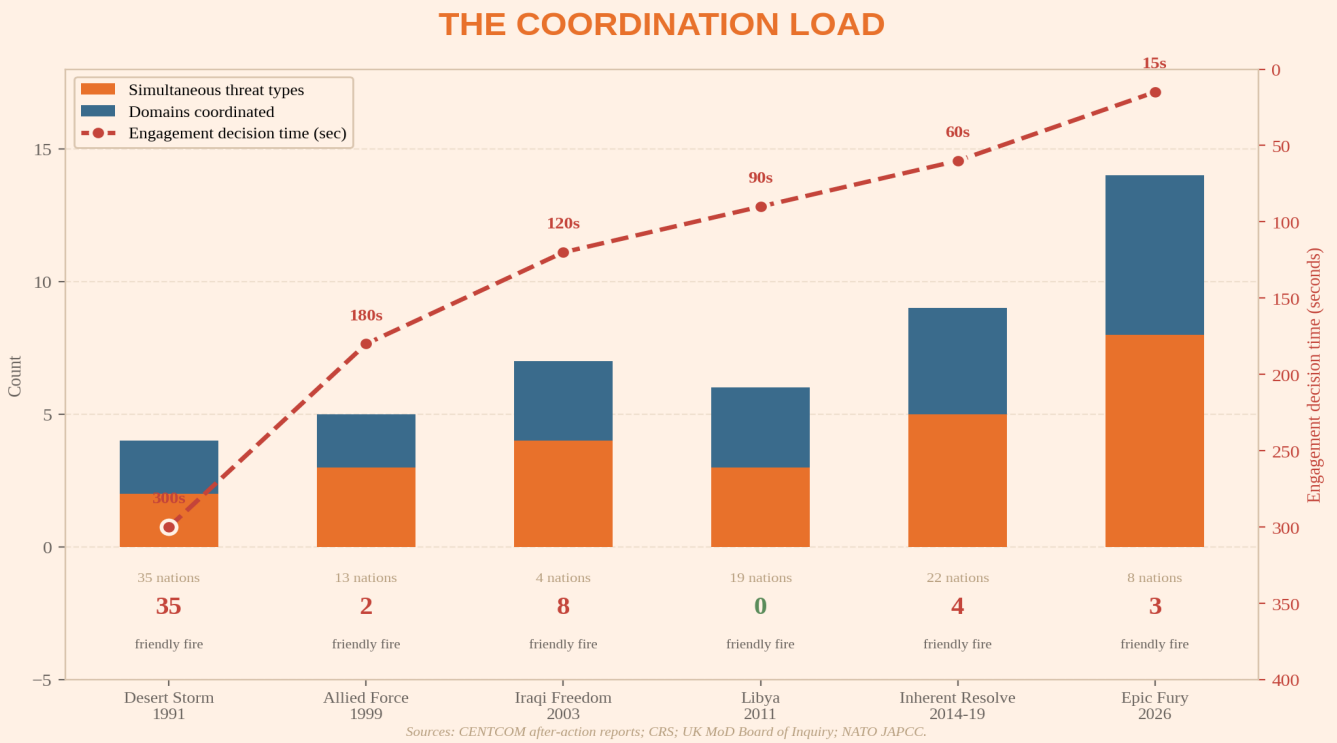

Figure 3. *The coordination load on coalition air defense architecture has increased monotonically across successive campaigns—more simultaneous threat types, more domains, less decision time—even as coalition size has varied. Friendly-fire incidents correlate with coordination load, not with technological sophistication or coalition size. Sources: CENTCOM after-action reports; CRS; UK MoD Board of Inquiry; NATO JAPCC.*

IMPLICATIONS FOR DEFENSE ECONOMICS

The spending fallacy. The conventional measure of military power is aggregate defense spending. Countries that spend more are presumed to be more capable, but multi-domain coordination costs are largely invisible in standard budget analysis as they are distributed across training budgets, communications programs, standardization offices, exercise schedules, and institutional overhead.

They do not appear as a discrete line item, and they are rarely measured against the coordination capacity they produce, making it possible to increase defense spending substantially while decreasing effective coordination capacity.

Hence, if new platforms are acquired without corresponding investment in the coordination architecture that integrates them, the result is more capability at the node level with lower performance at the system level. The military force becomes more expensive and more fragile simultaneously, which is the defense spending version of *architecture lag*: investment that accelerates technology without building the system around it.

The procurement trap. Current procurement systems evaluate platforms against domain-specific requirements. A fighter jet is assessed on its flight performance, sensor capability, and weapons capacity. These assessments are the military equivalent of Technology Readiness Levels: they measure component maturity. However they do not measure the platform's contribution to—or demand upon—the coordination architecture that makes multi-domain operations possible.

A platform that is individually excellent but architecturally expensive—requiring bespoke data links, non-standard identification protocols, unique doctrinal procedures, or dedicated coordination bandwidth—may reduce total system effectiveness even as it increases single-domain capability. Procurement systems that do not account for architecture cost will systematically over-invest in platform performance and under-invest in coordination capacity, resulting in an entire system-wide degradation.

The coalition penalty. Multi-domain operations are rarely conducted by a single nation. Coalition operations compound the coordination cost problem by introducing heterogeneous equipment, different national doctrine,

separate chains of command, and divergent rules of engagement. Every additional coalition partner multiplies the number of interfaces that must be managed and the institutional complexity of the coordination architecture.

The result is a structural penalty on coalitions: the more partners involved, the higher the coordination cost per unit of capability delivered. This penalty is invisible when operations are slow and permissive, but it becomes acute when the operational tempo demands decisions at machine speed across national boundaries. The architecture that works in peacetime exercises, where tempo is controlled and ambiguity is scripted, may fragment under the pace and uncertainty of contested multi-domain combat.

This has direct implications for alliance design, in that an alliance optimized for maximum membership may be structurally worse at multi-domain coordination than a smaller, more deeply integrated force.

For example, NATO's expansion of membership increases its aggregate capability while potentially decreasing its effective coordination bandwidth, particularly in scenarios where interoperability has not kept pace with political inclusion.

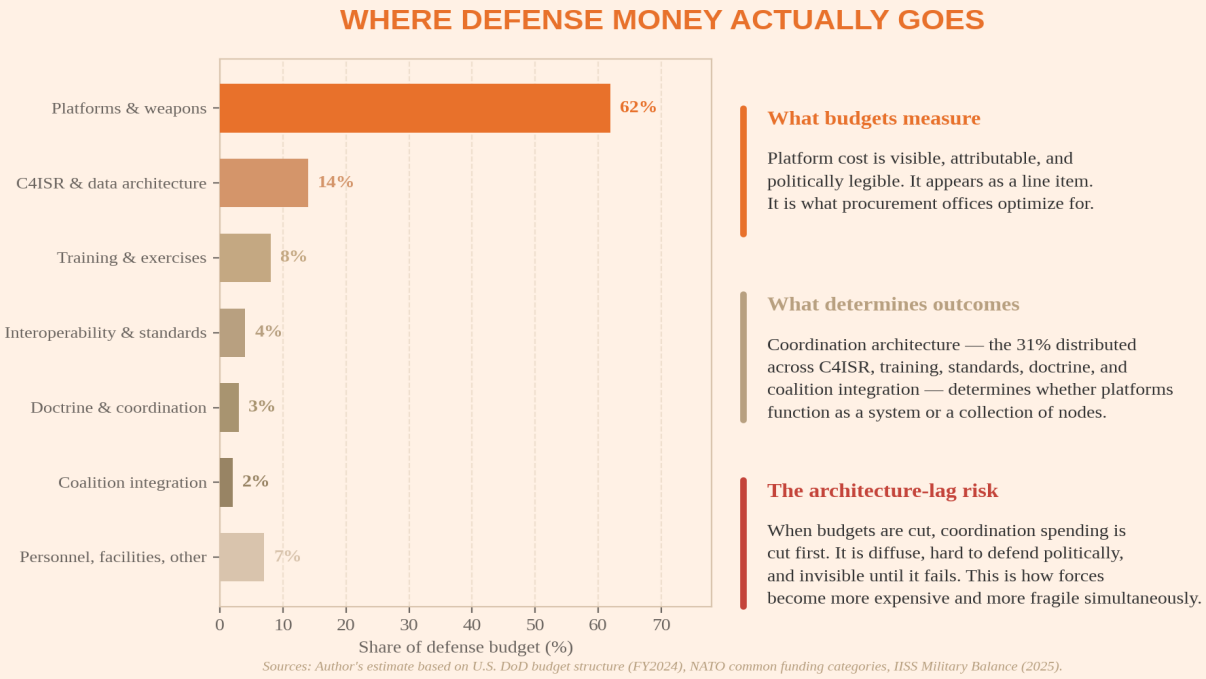

Figure 4. *Platforms and weapons account for roughly 62% of defense budgets. The remaining 38%—the coordination architecture that makes platforms interoperable—is distributed across budget lines rarely measured as a system. This is the expenditure that determines whether multi-domain integration works or fragments. Sources: Author's estimate based on U.S. DoD budget structure (FY2024), NATO common funding categories, IISS Military Balance (2025).*

The integrating force must synchronize every interface, but the disruptor must only find one that fails. The economics of this exchange are structurally unfavorable to the integrator—and they worsen with every domain added.

THE STRATEGIC CALCULUS

None of this argues against multi-domain integration, of course. A force that can coordinate across domains retains enormous advantages in situational awareness, operational flexibility, and the ability to concentrate effects from multiple vectors. Instead, the argument is that these advantages come at an architectural cost that is poorly understood, rarely measured, and systematically underestimated—and that this cost creates structural

vulnerabilities that asymmetric adversaries are rationally incentivized to exploit.

The strategic calculus for the integrating force is therefore not simply “more domains, more advantage.”

It is a *constrained optimization*: how many domains can the coordination architecture actually synchronize at the speed and reliability required, given the threat environment it must operate in? Adding a domain whose coordination cost exceeds the capability it contributes makes the force more expensive and more vulnerable simultaneously.

The optimal multi-domain force is not the one that integrates the most domains, it is the one whose coordination architecture can sustain integration at the bandwidth the operational environment demands.

For the asymmetric adversary, the calculus is the mirror image and the optimal strategy is not to compete within all domains, but to identify the interfaces between domains where the coordination architecture is weakest and impose costs there.

The economics favor this strategy overwhelmingly: A GPS jammer costs orders of magnitude less than the timing synchronization architecture it disrupts, and a drone swarm costs a fraction of the air defense coordination bandwidth it consumes.

The result is a structural dynamic in which the integrating force’s spending on coordination architecture is disciplined by the disruptor’s ability to target interfaces cheaply. The more the integrator spends on architecture, the more interfaces exist for the disruptor to target; and the more the disruptor targets interfaces, the more the integrator must spend on resilience. This is not a stable equilibrium, it is an arms race in which the integrator’s costs grow faster than the disruptor’s.

The countries and alliances that will sustain military advantage in this environment will not be those that build the most domains into their force structure, as is the current stated strategic goal of the United States. They will be those that build coordination architectures with sufficient bandwidth to process multi-domain complexity at the speed the threat demands, at a cost that does not bankrupt the system it is meant to protect.

Multi-domain superiority is real, but it is not a *platform* property, it is an *architecture* property. But unfortunately, architecture has a cost curve that the current defense-economic framework does not adequately measure, manage, or even see.

This article draws on frameworks developed in two working papers by Sinéad O’Sullivan: Institutions as Coordination Architectures: Adaptive Bandwidth and the Dynamics of Economic Development and Market Formation as a Systems Engineering Problem. Both papers develop the formal models, cross-domain evidence, and theoretical infrastructure for the coordination cost and architecture lag frameworks applied here to military operations.

Copies are available on request at s@sinead.co.