

# The Chips Race Is a Systems Race

Why Architecture, Not Technology, Will Decide the US-China Semiconductor Competition

Sinéad O'Sullivan

## IDEA IN BRIEF

### THE PROBLEM

US semiconductor strategy treats the China competition as a fabrication race—whoever builds at the most advanced nodes wins. Export controls and fab subsidies reflect this ideology. But semiconductor leadership has never been determined by fabrication alone, and controlling one layer of the stack has never been sufficient to sustain industrial dominance.

### THE ARGUMENT

The real contest is over coordination architecture: the ability to synchronize technical capability with supply chains, standards, procurement, workforce systems, and market demand. China is building an architecture-first strategy—mature-node capacity, domestic equipment ecosystems, and AI inference deployment systems designed to turn good-enough chips into functioning industries.

### THE IMPLICATION

As AI deployment shifts from training to inference at scale, the advantage moves from frontier chips to deployment systems. The risk for the US is not that China produces a superior GPU but that China constructs a superior system around an inferior one.

**T**he story the US tells itself about the

semiconductor competition goes like this: because we control the chokepoints through export controls on advanced lithography, restrictions on AI chip sales, and alliances with Dutch and Japanese toolmakers, Washington has built a technological blockade designed to keep China a generation behind in chip fabrication. If we control the most advanced nodes—3nm, 2nm—we can control the future of AI, defense computing, and economic power.

The assumption underneath all of this is that the semiconductor race is a technology race, in that whoever fabricates the most advanced chip wins, because the ecosystem will follow.

But it won't.

Decades of evidence from development economics and systems engineering point to a broader conclusion that technologies don't produce markets on their own. Rather, markets emerge when

technical capability, market structures, and institutional coordination evolve in synchrony. And on that dimension—the alignment of the full system, not just the leading component—the competitive picture is considerably more contested than the export-control narrative suggests.

## THE ARCHITECTURE GAP, NOT THE TRANSISTOR GAP

To see why, it helps to distinguish between two kinds of advantage.

The first is **component-level advantage**: superiority in a specific technical subsystem. In semiconductors, that means fabricating chips at the smallest process nodes. TSMC, Samsung, and Intel compete here, and the US-led coalition has a clear lead.

The second is **architecture-level advantage**: the ability to align technical capability with the supply chains, standards, procurement systems, workforce pipelines, and demand structures that turn chips into industries. This is the harder, less visible dimension

of competition. And it's where the real contest is playing out.

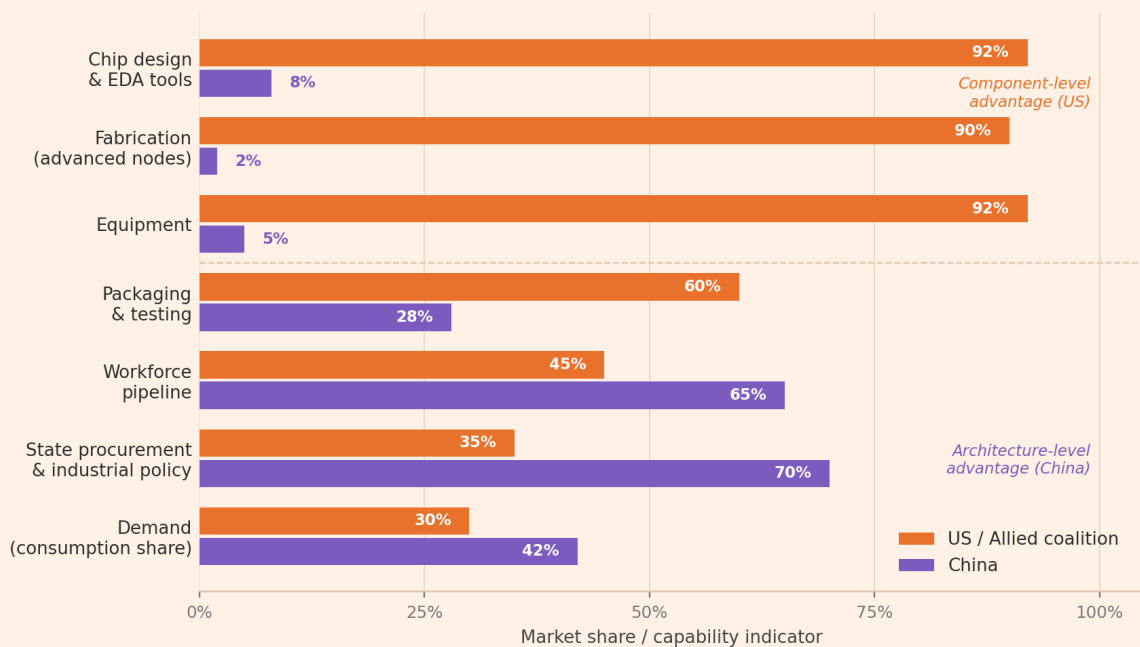
Semiconductor leadership has never been determined by fabrication alone. Japan's rise in the 1980s wasn't just DRAM fabrication—it was the coordination of equipment suppliers, materials science, and industrial customers into an integrated production system. Taiwan's emergence as the world's foundry rested on TSMC's business model, but also on decades of government coordination

across education, infrastructure, standards, and procurement.

Similarly, South Korea's success depended on Samsung and SK Hynix, but equally on the institutional architecture—directed credit, export coordination, workforce development—that allowed those firms to scale.

In every case, the winning move wasn't building the best chip in isolation, it was building the system around the chip.

### THE SEMICONDUCTOR ARCHITECTURE STACK



Sources: SIA Factbook 2025; CSIS Mapping the Semiconductor Supply Chain (2024); TrendForce OSAT Rankings (2025); SEMI World Fab Forecast (2024); SIA/BCG Workforce Study (2023). 'US/Allied' includes US, Taiwan, Japan, S. Korea, Netherlands, EU.

**Figure 1:** Market share across the semiconductor architecture stack. The US-led coalition dominates the top layers—EDA tools, equipment, advanced fabrication—while China leads on demand, state procurement, and workforce scale. The contested middle (packaging and testing) is where the balance is shifting fastest. Sources: SIA Factbook 2025; CSIS (2024); TrendForce (2025); SEMI (2024); SIA/BCG (2023).

*The risk for the US isn't that China produces a superior GPU. It's that China constructs a superior system around an inferior one.*

## WHAT CHINA IS ACTUALLY DOING

This reframes China's semiconductor strategy entirely, because what looks like a failed attempt to replicate TSMC could be better understood as a deliberate exercise in architecture-building.

Yes, SMIC's most advanced processes lag the frontier by several nodes. But the investments being made across China's semiconductor ecosystem aren't concentrated on closing that transistor gap, they're distributed across the full stack of what it takes to turn chips into functioning economic and military systems.

Consider what's being built simultaneously: massive expansion of trailing-node and mature-node fabrication—the chips that actually go into cars, factories, IoT devices, and consumer electronics; a

domestic equipment ecosystem (SMEE, Naura, AMEC), still behind ASML and Tokyo Electron, but increasingly capable at nodes sufficient for most commercial and defense use; an AI deployment architecture that prioritizes inference at scale over training at the frontier—a strategy that needs enormous volumes of capable-enough chips rather than small quantities of the best ones; and procurement and standards systems designed to absorb domestic production, creating the demand-side architecture that turns manufacturing capacity into market viability.

This is the deliberate pursuit of an *architecture-first* strategy. The question it poses to Western policymakers is dizzying: *can a synchronized system built around trailing-node chips outperform a fragmented one built around leading-node chips?*

## TWO STRATEGIES COMPARED

DIMENSION	US / ALLIED APPROACH	CHINA'S APPROACH
Primary lever	Export controls & fab subsidies targeting frontier nodes (3nm, 2nm)	Mature-node capacity expansion across automotive, IoT, industrial, consumer
Equipment strategy	Chokepoint control via ASML, Tokyo Electron, Applied Materials restrictions	Domestic equipment buildout (SMEE, Naura, AMEC) for trailing nodes
AI hardware focus	Frontier training accelerators (NVIDIA H100/B200 via TSMC)	Inference at scale (Huawei Ascend ecosystem with integrated software & procurement)
Demand architecture	Market-driven; fragmented procurement across agencies and firms	State-coordinated procurement guaranteeing domestic absorption
Institutional speed	Optimization architecture: slow regulation, competing standards bodies, risk-averse procurement	Transformation architecture: rapid policy iteration, centralized coordination, high adaptive bandwidth

## THE WESTERN VULNERABILITY

The US and its allies face a structural challenge that export controls can't resolve. Western semiconductor policy is overwhelmingly focused on the supply side: subsidizing fabrication through the CHIPS Act, restricting technology transfer, investing in frontier R&D. These are important of course, but they address only one layer of the architecture.

The deeper problem is that Western institutional systems are optimized for stability, not speed: Regulatory processes are slow and fragmented; standards development is distributed across

competing bodies with limited coordination; procurement is designed to minimize risk and ensure procedural compliance, not to create early demand for strategic technologies; infrastructure provision—from power grids to advanced packaging facilities—is politically contested and chronically delayed.

These aren't policy failures. Rather, they're structural features of *optimization architectures*—institutional systems built to manage mature economies with predictability and low variance. They're excellent at maintaining existing industries, but they're also poorly suited to building new ones at speed.

The result is a growing mismatch between investment and readiness. The US is directing tens of billions toward fab capacity while the coordination architecture needed to make that capacity strategically effective—workforce pipelines, packaging ecosystems, demand aggregation, supply-chain redundancy, regulatory streamlining—develops far more slowly. The CHIPS Act funds factories, but factories without synchronized architectures can only produce chips that will struggle to find their way into the systems that count.

### THE AI HARDWARE INFLECTION

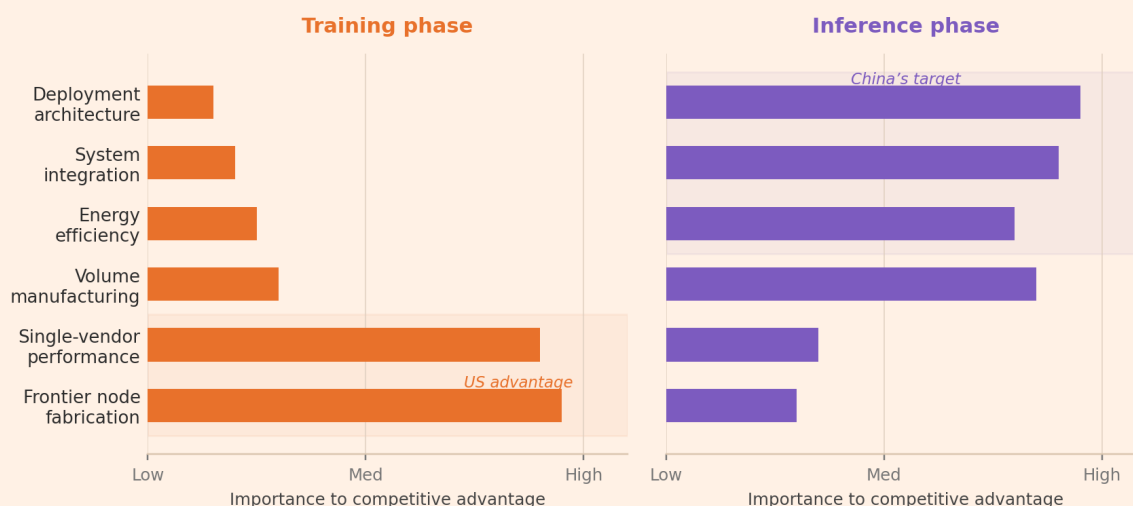
This dynamic is sharpest in AI hardware, where the architecture question is becoming urgent. The current US advantage rests on a narrow foundation: NVIDIA’s dominance in training accelerators, manufactured by TSMC in Taiwan, using ASML equipment from the Netherlands. Potent—but

geographically concentrated, single-vendor-dependent, and disruption-vulnerable.

The next phase of AI shifts the importance from training to inference—from building models to running them at scale across every sector of the economy. Inference is less dependent on frontier nodes. It rewards volume, energy efficiency, system integration, and the ability to deploy hardware into diverse operational environments. In short, it’s an architecture problem, not a transistor problem.

China’s strategy is explicitly oriented toward this transition. Huawei’s Ascend AI chips, while less performant than NVIDIA’s best, are designed to function within an integrated ecosystem of software frameworks, cloud infrastructure, and government procurement that guarantees demand at scale. The architecture is being built around the chip, not the other way around.

### THE TRAINING-TO-INFERENCE SHIFT



**Figure 2:** In the training phase, competitive advantage concentrates on frontier fabrication and single-vendor performance—where the US leads. In the inference phase, it shifts to volume, efficiency, integration, and deployment architecture—where China is building.

### COMPETING ON COORDINATION

None of this means the US has lost the semiconductor competition, as American firms still lead in chip design, EDA tools, and frontier

fabrication partnerships. The allied coalition retains formidable advantages in equipment, IP, and talent.

But the terms of the competition are shifting in a direction that Western policy frameworks are poorly equipped to address. The countries that built world-leading semiconductor industries—Japan in the

1980s, Taiwan and South Korea from the 1990s onward—didn't succeed by excelling at one layer of the stack. History shows us that they succeeded by acting as architects of synchrony: aligning standards, infrastructure, finance, procurement, workforce development, and industrial demand into coherent systems.

Export controls buy time in the same way that fabrication subsidies build capacity, but neither creates the coordination architecture that converts this capacity into a sustainable strategic advantage.

Winning the semiconductor race requires something harder than restricting access to lithography machines or directing capital toward new fabs. It requires building the institutional machinery—fast, adaptive, and coordinated—that can synchronize technical investment with market formation, workforce readiness, and deployment at scale.

The chip race is a systems race. And systems races are won by whoever builds the architecture first.

---

*This article draws on 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, mathematics, and cross-domain evidence summarized here.*

Copies are available on request at [s@sinead.co](mailto:s@sinead.co).