DARPA Will Decide the Fate of Your Quantum Investment!

ICM HPQC FUND

Matthew Gould
Portfolio Manager, ICM HPQC Fund

December 2025

The Utility Race: How DARPA’s Benchmarking Initiative Could Decide the Future of Quantum Computing

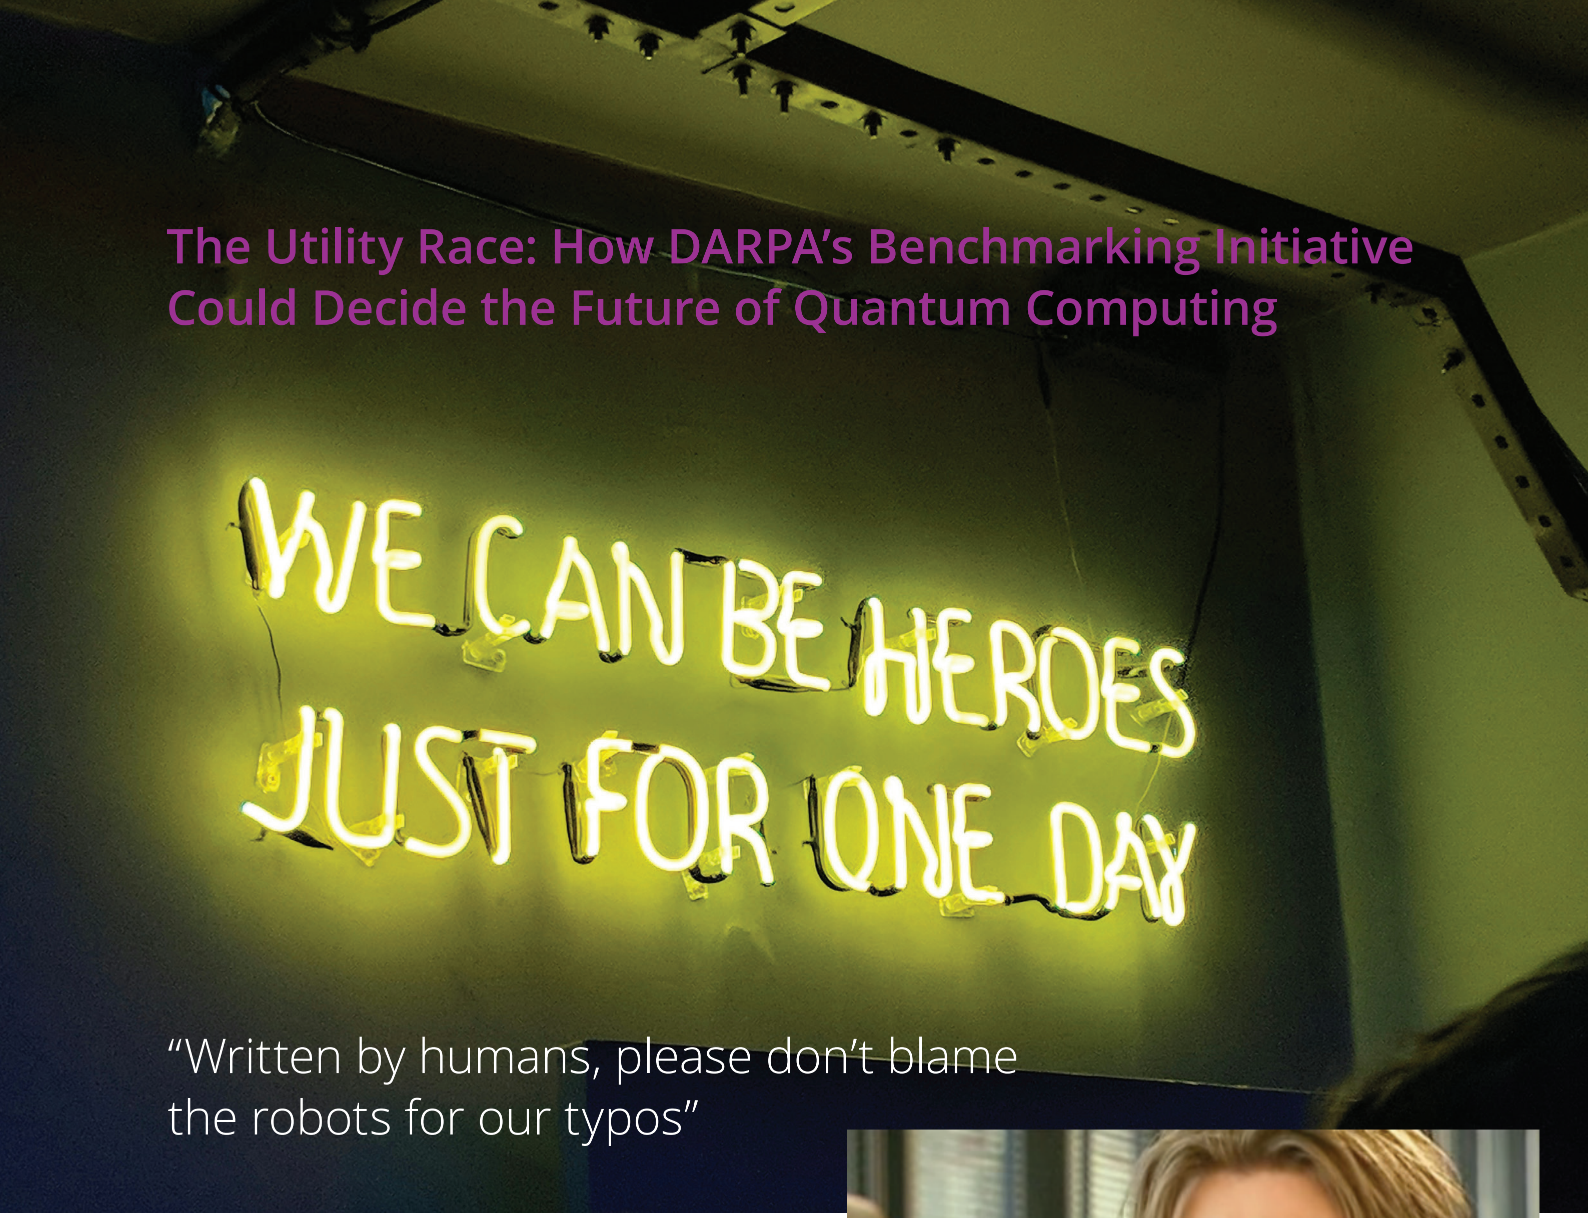

“Written by humans, please don’t blame the robots for our typos”

Introduction: The Reckoning

Back in 1999 probably the coolest man to ever walk the earth, David Bowie, sat down with the historically uncool Jeremy Paxman and gave a remarkably prescient view of how this new thing called the ‘internet’ was going to impact life, music and media.

You can see it **here** (skip forward to the 3.20 min mark if you want to get straight to the forecast). Bowie describes his sense that a profound change was about to sweep across most aspects of culture and business due to the internet.

On any lookback Bowie was right – the internet of 1999 was just about to break its bounds, and, once through the rationalising sense-check of the dotcom bust changed almost every aspect of society. Big vision statements by new age tech gurus are now commonplace and even form the centrepiece of sober analysts' reports on things like AI and robotics. None of the "tech bros" are remotely as cool as Bowie. So accustomed are we to vision peddling that some VC's bank on it – planning their investment strategies to catch the 'hype-wave' – getting into a sector on the rise, and then out at the peak – usually just before commercial reality hits the projects they invested in. Nowhere is this more the case than in the field of quantum computing.

Quantum computing has been sold as the coming industrial revolution (it will be). Consultants speak in trillions; start-ups promise miracles in chemistry, finance, logistics and defence. Yet beneath the marketing thrives an awkward worry: the most prominent quantum machines being built today – by companies such as IonQ, PsiQuantum, Rigetti and IBM – may never be commercially viable. They are costly to build, more costly to operate, and currently incapable of solving any problem whose answer is worth that cost.

A new American government programme is quietly forcing the industry to face this reality. DARPA's Quantum Benchmarking Initiative (QBI) is a forensic audit: technically deep, commercially brutal and modality-agnostic. Its mandate is unforgiving – **determine which quantum approaches, if any, can deliver utility-scale quantum computers within a decade**. A utility-scale quantum computer is one whose economic value exceeds its capital and operating costs.

For investors, strategists and policymakers, **QBI is the most consequential quantum program** in the world.

Cost: The quantum compute problem that won't go away…

Despite the charm of laboratory demonstrations and sleek marketing decks, the economics of today's quantum hardware are dire. To deliver valuable results – in catalysis, drug discovery, cryptography or materials science – a quantum computer requires thousands of **logical qubits**, each protected by elaborate error-correction codes built from millions of **physical qubits**. None of today's machines are even close. The field's frontrunners have only a few dozen or a few hundred noisy physical qubits.

Scaling current architectures is a nightmare. Superconducting systems require ever-larger cryogenic enclosures. Ion traps demand extra lasers, optics, vacuum chambers and control electronics. Photonic systems require blizzards of photon sources and detectors. All require escalating amounts of power, space and capital.

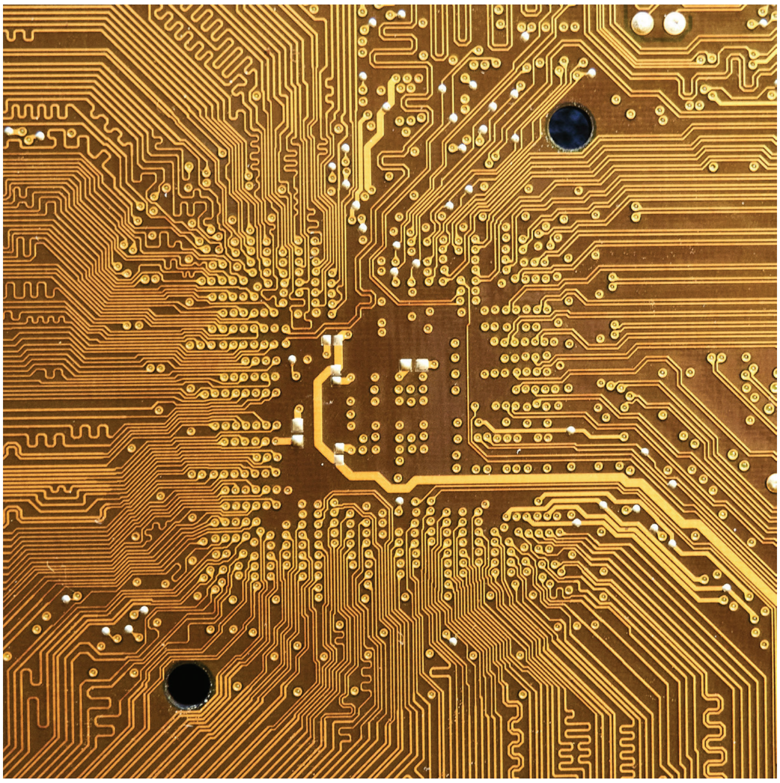

Worse, when scaled to the sizes needed for meaningful work, many architectures exceed the physical limits of data centres. Some would demand more than 100 megawatts of power – more than the world's largest supercomputers. Others require room-sized vacuum apparatus or thousands of modular connections that render mass production absurd.

It is little wonder that the financial performance of quantum start-ups has drifted away from expectations. Revenues remain modest and are far exceeded by costs; losses mount; investor confidence yo-yos with each scientific press release or tech-CEO plaudit. The sector has long suffered from the convenient illusion that commercial feasibility can be postponed indefinitely. QBI ends that illusion.

DARPA's definition of utility

DARPA's definition of a **utility-scale quantum computer** is disarmingly simple: a system whose **computational value exceeds the full cost of producing it** – including capital expenditure, energy, maintenance, personnel and integration.

This definition is less about physics than economics. It asks whether a data-centre operator, national lab or industrial user would gain more value from the machine than they spend to run it. It reframes quantum computing as a business proposition rather than a science experiment.

Under this definition, the focus turns to which architectures can plausibly:

- Achieve **thousands** of logical qubits in compact systems
- Deliver multi-day fault-tolerant computation
- Fit inside existing data-centre footprints
- Operate within reasonable power limits
- Be built by the tens or hundreds, not as bespoke prototypes
- Support commercially relevant workloads at commercially acceptable costs

It is a definition that rewards manufacturability and punishes complexity.

QBI is a three-stage programme, each stage escalating scrutiny, detail and funding. Its structure is designed to mimic the due diligence process of a highly sceptical investor – except with hundreds of physicists, engineers, coders and economists involved.

Stage A – Plausibility (6 months, up to $1m per participant)

Firms propose a **complete utility-scale concept**, not a prototype. They must present:

- A full architectural blueprint
- Fabrication strategy
- Control electronics
- Error-correction plan
- Power, cost and scaling models
- A realistic path to thousands of logical qubits

An independent **test and evaluation (T&E)** cadre – spread across national labs, universities and specialist contractors – rips each plan apart. Stage A ends with a winnowing: only credible approaches survive. This is the stage that has just been completed.

Stage B – Detailed R&D plan (12 months, up to $15m per participant)

Survivors receive significantly more funding. In Stage B they must deliver:

- Detailed engineering and fabrication roadmaps
- Subsystem prototypes
- Validated error-rate improvement plans
- Risk mitigation analysis
- Verifiable cost curves
- Demonstrated subsystem performance

This is the stage at which grand narratives will tend to die. Companies must show not only how to build the pieces but how to assemble them into a machine whose **economics** make sense. Some receive additional co-funding from national governments or strategic partners.

Stage C – Validation and co-design (tailored duration, up to $300m)

The very few entrants that clear Stages A and B may receive large-scale funding – **up to $300m each** – to build and demonstrate actual hardware aligned with their utility-scale goals. Independent evaluators run real workloads on these machines to verify claims about performance, power consumption and scalability.

QBI's funding is structured like a funnel. Stage A spreads small cheques widely to test ideas; Stage C concentrates far larger cheques on approaches that have earned credibility.

What DARPA is benchmarking

DARPA is not benchmarking toy problems. Many of the claims of 'commercial problem solving' in today's quantum computing sector relate to tightly controlled 'sandpit challenges' – problems designed as much to suit the specific hardware each company has built as they are to prove a single limited commercial 'use-case'. DARPA meanwhile has selected a set of application-level workloads that (a) have significant economic or national-security value and (b) require capabilities far beyond today's machines.

The benchmark portfolio includes:

- **Catalyst and battery chemistry**
- **High-value materials simulation**
- **Corrosion-resistant alloys**
- **Metalloprotein binding and drug-design pathways**
- **Fermi–Hubbard and strongly correlated materials**
- **Magnetic and high-pressure materials modelling**
- **Rocket propellant discovery**
- **Cryptography workloads, including RSA-2048**

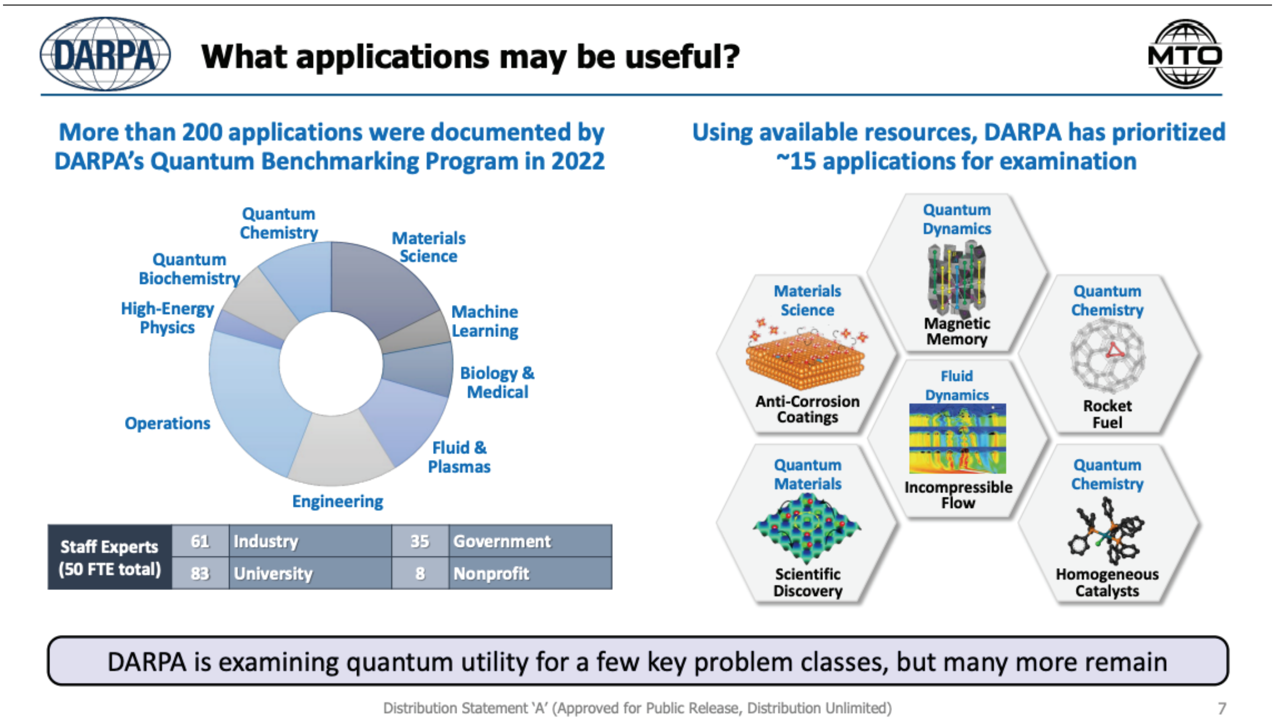

For each of these workloads, DARPA's teams and programme participants estimate:

- Required logical-qubit count
- Total logical operations
- Surface-code or other error-correction overhead
- Physical-qubit requirements under each modality
- Expected wall-clock runtime
- Hardware footprint, energy budget and cost profile

These benchmarks act as a corrective to hype. They show, for example, that many attractive chemistry problems require thousands of logical qubits – far beyond the limits of machines that occupy entire rooms today or even the limits of machines targeted in the future. Here's what this chart shows – it's worth taking this one in:

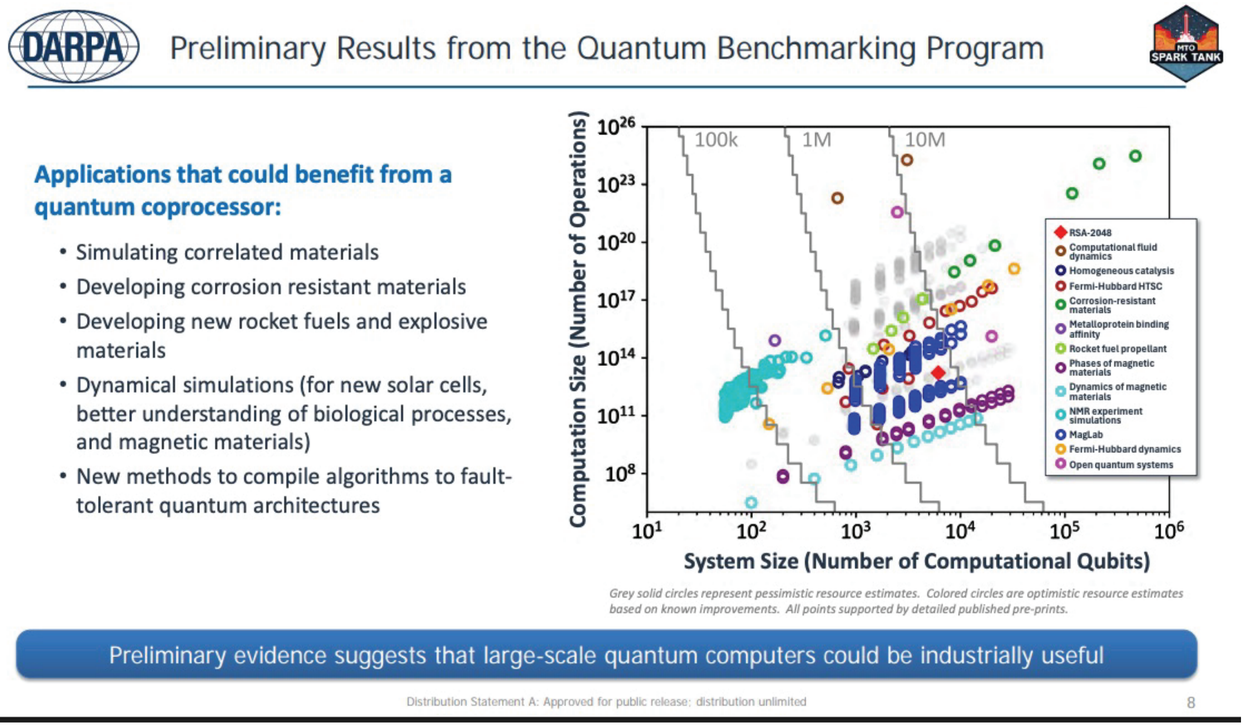

The grey lines in this graph refer to the number of **physical qubits** required for the number of logical qubits (called computational qubits in this graph). The coloured points represent the estimated resource requirements for specific applications as researched by DARPA. Notice that most resource estimates require a quantum computer with greater than 1M physical qubits, with some applications exceeding 10M physical qubits. It reinforces that scalability is a critical metric for quantum computing companies and that a 1M physical qubit quantum computer is not a target – it's table stakes.

This chart is a useful reference as quantum companies continue to scale qubit numbers, with some expecting devices with thousands of qubits in 2026. Companies with high qubit numbers might appear to lead the pack, but this chart is a sober reminder that commercial success requires scaling into the millions of qubits regime.

QBI participants will need to convince DARPA that they can scale to 1M physical qubits and eventually beyond this number to progress through all three stages of the programme.

Shor's algorithm: still the canonical stress test

Shor's algorithm for factoring large integers remains the most famous quantum workload. Factoring a 2,048-bit RSA key is a canonical milestone because it sits near the upper end of quantum capability and near the lower end of national-security urgency. If a quantum computer can crack standard messaging and financial transaction encryption your government really needs to own one. While Shor's algorithm appears in almost every quantum computing pitch deck, or McKinsey report on the TAM of the quantum market, few analysts go deeper and ask just what it will take from a compute power perspective to deliver results from running this code at commercial scale. Things get informative when you do.

DARPA others suggest that factoring RSA-2048 with full error correction requires:

- A few thousand **logical qubits, which would, depending on the modality of qubit being used require**
- Between **one million and twenty million physical qubits** (depending on error rates, gate speeds and code cycles)
- Hours to days of runtime
- Tight control of noise and decoherence throughout

Machines that cannot support workloads of this scale cannot be considered "utility-scale" in DARPA's sense. The point is not to crack encryption; it is to establish whether an architecture can support the largest, most resource-intensive jobs foreseen in the next decade. Then do it repeatedly, and for multiple customers at one.

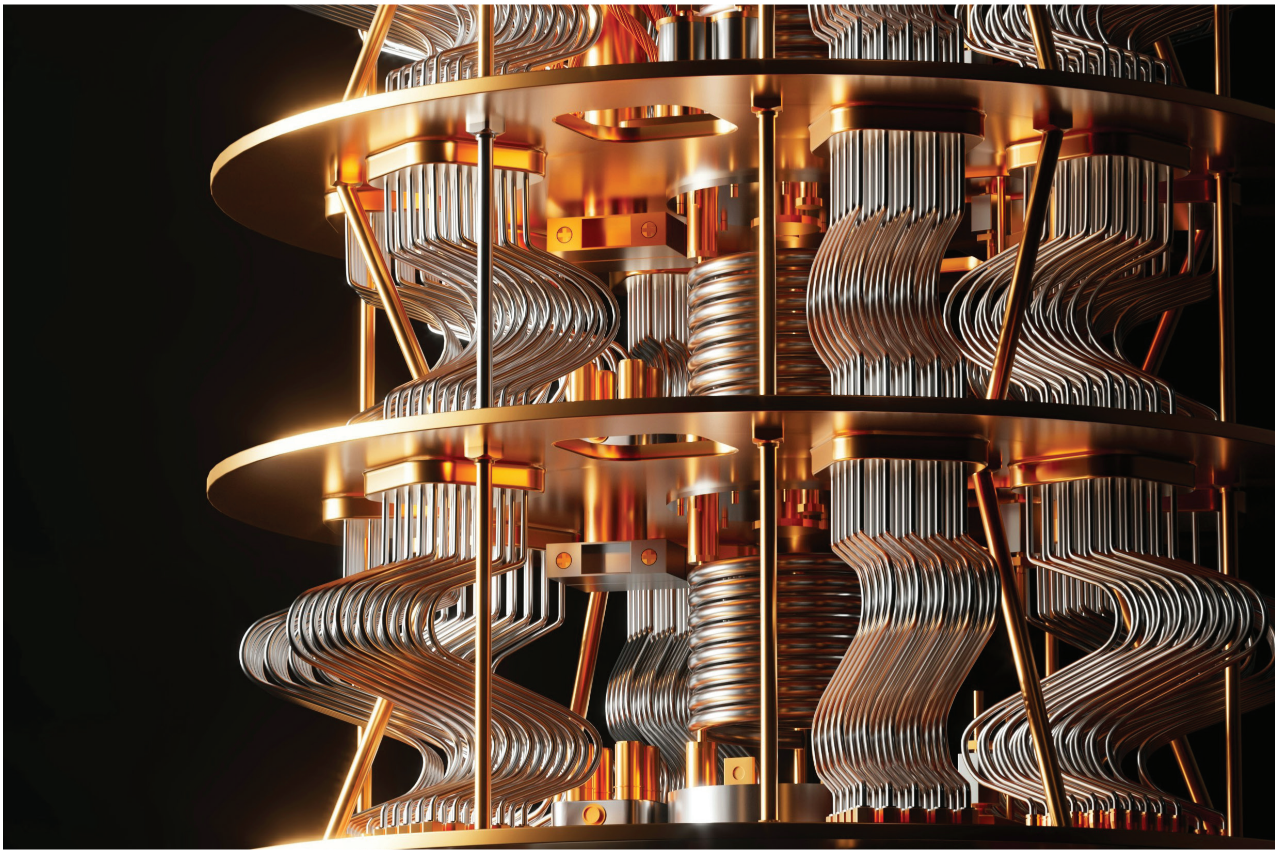

The data-centre test

At the HPQC fund we loves us a Data Centre. It's what we do do...

We specialise in selecting the compute technology that is currently early-stage but will be dragged into production within the next five years as the pace of change at a DC level continues to increase. So we know, I mean we really KNOW that, for a quantum computer to be commercially viable, it must behave like any other high-value data centre asset. That means:

- **Footprint:** The system should fit in a few square metres, not a custom-built facility.
- **Power:** A machine should ideally consume **under 10 MW**, and certainly under 40 MW.
- **Manufacturability:** It must be possible to build dozens, hundreds or even thousands of them economically and position them within standard high performance data centres globally.
- **Integration:** They must co-exist with classical hardware, not require bespoke infrastructures.
- **Cooling:** Cryogenic load must be manageable with modern systems.

When assessed under these criteria, four of the five dominant modalities struggle. Their power draw, physical size and control complexity grow unmanageably as qubit counts scale.

Spins in silicon: the disruptive contender

The quiet outlier – **spins in silicon** – is beginning to look like the most commercially plausible route to utility-scale quantum computing.

Its advantages are mundane but decisive:

- Spin qubits can be fabricated at sizes similar to advanced transistors. In principle, **billions of qubits** could be integrated on a single chip
- Unlike other modalities, silicon spin qubits can be fabricated in existing semiconductor foundries using standard processes. This allows quantum processors to piggy-back on trillion-dollar fabrication ecosystems and enjoy Moore's-law-style scaling. Spins in silicon don't require its own industrial revolution, its leveraging a revolution that's already happened. (Ref. John Martnis FT article)
- A utility-scale silicon-spin machine could fit in **one cryostat**, not hundreds of interconnected modules. Its footprint could be similar to that of a refrigerator-sized HPC node.
- Because much of its architecture can use advanced but familiar CMOS electronics, the expected energy budget is **two orders of magnitude lower** than many competitors.

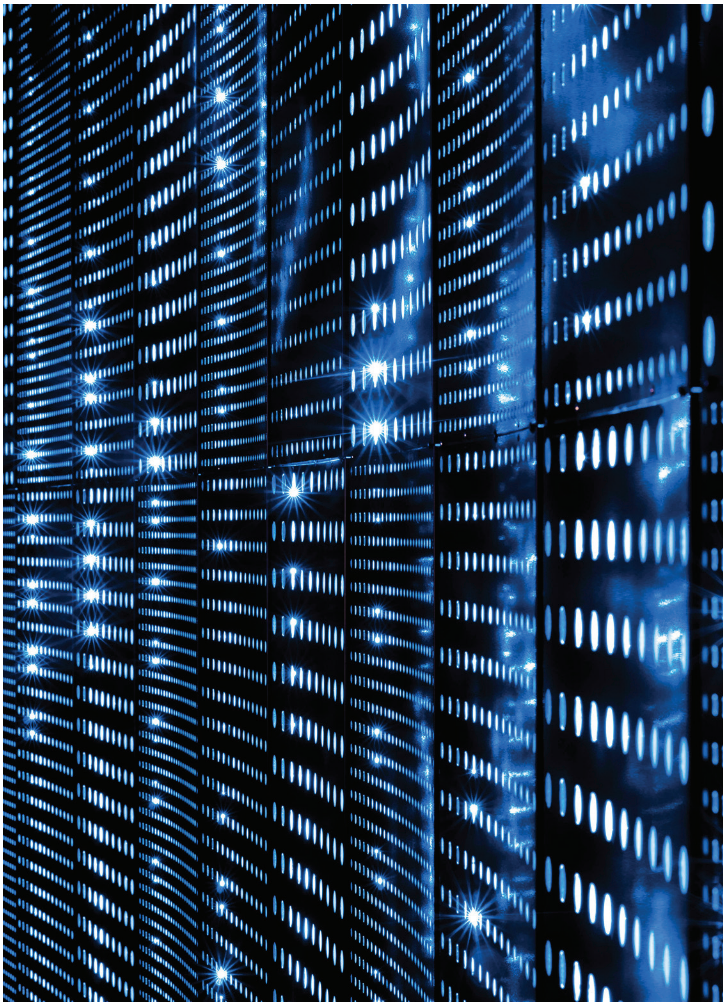

This combination has not gone unnoticed. Silicon-spin firms such as Diraq and Quantum Motion are among those advancing through QBI's selection rounds. **If they deliver what their roadmaps promise, they could leapfrog more established modalities** in terms of qubit numbers and offer a significant pricing advantage

A timeline to 2033

DARPA's timeline is explicit. It expects:

2024 to 2026	2025 to 2027	2027 to 2030	By ~2033
Broad benchmarking of all modalities	Stage-A and Stage-B down-selection	Stage-C hardware evaluations	Determination of whether any architecture can deliver at **least one industrially valuable utility regime**

Industry is already responding. PsiQuantum is building a massive quantum facility in Chicago that will serve, in part, as a platform for QBI's assessments. Others are preparing their own hardware for Stage-C consideration.

The **next eight years will define** which quantum technologies survive and which fade into historical curiosity.

QBI is, in effect, the world's most expensive and sophisticated technical due-diligence machine. For investors, it matters for three reasons:

1. It reveals which technologies are structurally viable. QBI's questions are unforgiving. Architectures with bad physics will be cut out. Those with elegant physics but poor manufacturability will not progress.
2. It rewards alignment with semiconductor economics. The entire semiconductor industry has spent more than half a century learning to build nanoscopic devices cheaply and at scale. QBI implicitly favours modalities—like spins in silicon—that can harness that ecosystem.
3. It compresses the time horizon. Instead of vague promises of "quantum advantage in ten years", QBI forces companies to show near-term, verifiable progress.

Investors who ignore QBI do so at their peril.

The coming realignment

If QBI confirms that the majority of today's architectures cannot scale to utility levels within the decade, the quantum industry will undergo a brutal correction. Capital will migrate away from designs that are physics-rich but economically doomed.

But if even one technology – especially a CMOS-compatible one – passes the Stage-C tests, the rewards could be spectacular. The winning firm could become the Intel or Nvidia of the quantum era, supplying machines that sit at the heart of scientific discovery, industrial optimisation and national-security infrastructure.

The next phase of QBI may thus be remembered as the moment quantum computing stopped being science fiction and became a capital-intensive industry ruled not by qubit counts, but by manufacturability, economics and utility.

Or, in the words of Sir David Bowie "Tomorrow belongs to those who can hear it coming".

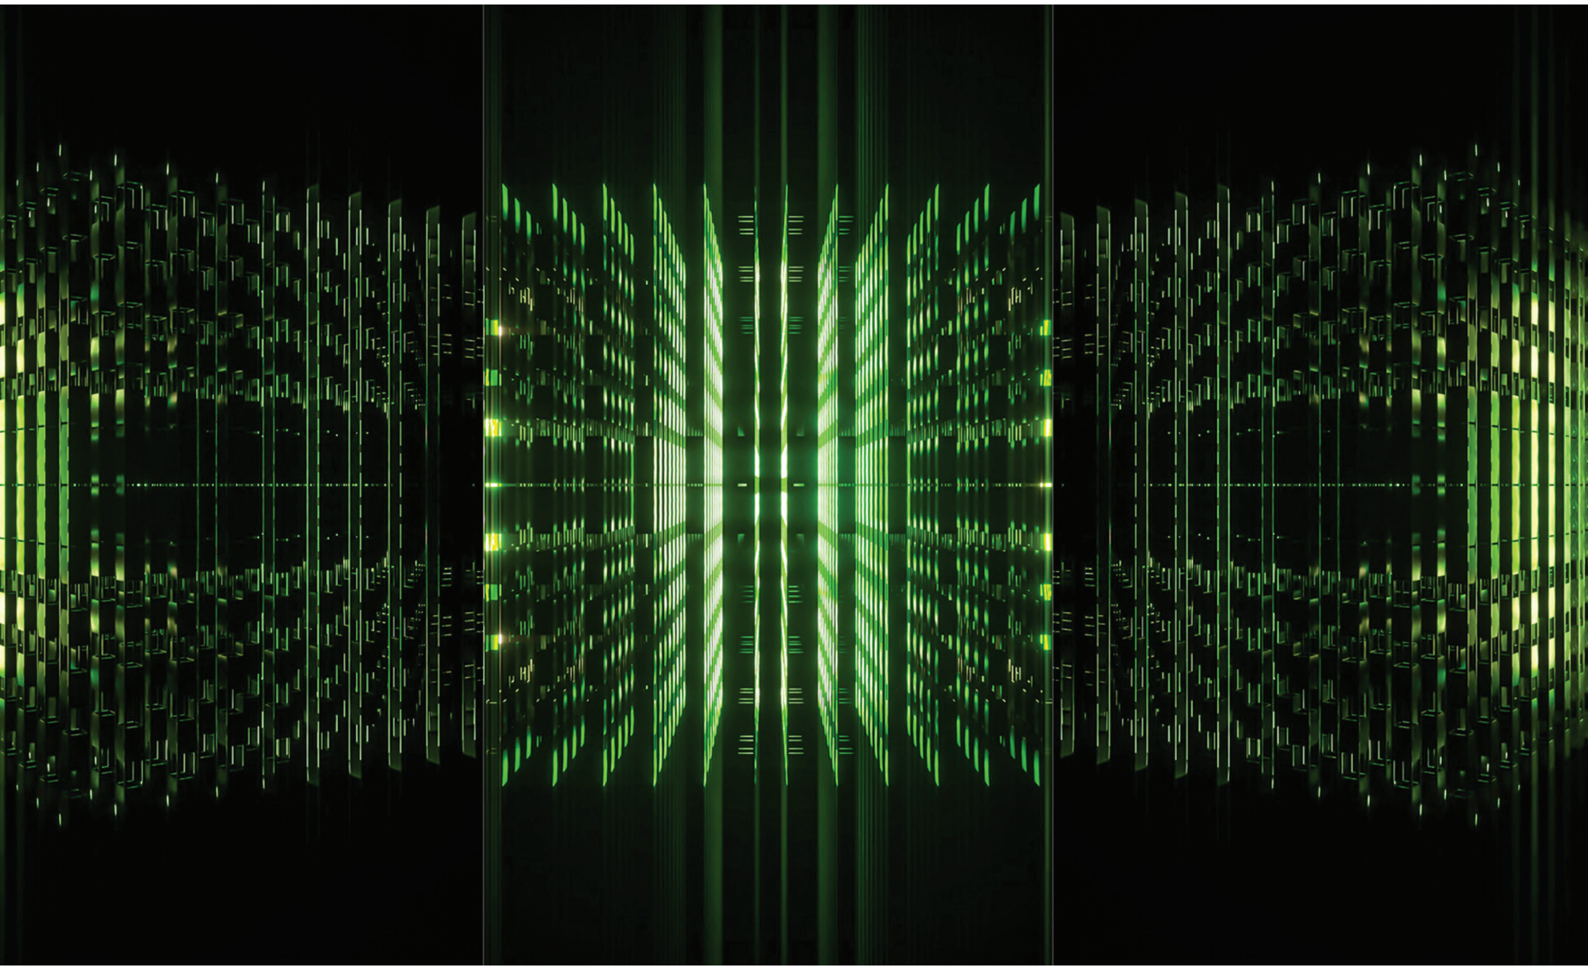

Disclaimer