

Beyond Silicon, Beyond the Qubit

Quantum Biology, Graphene Photonics, and the Competing Architectures of Future Computation

Stephen Bray: Family Business Consultant

Synopsis

This paper argues that the future of computation will emerge from the convergence of four domains: quantum biology, graphene photonics, superconducting quantum error correction, and competing device architectures for scalable quantum computing. Quantum biology has matured into a recognised research field concerned with non-trivial quantum effects in living systems, including photosynthesis, enzyme dynamics, and magnetoreception (Cao et al., 2020). Graphene photonics has developed as a promising response to the energy and bandwidth limits of copper interconnects, with graphene's broadband optical response and electrical tunability making it a candidate material for compact modulators and photonic integration (Bonaccorso et al., 2010). In quantum hardware, Google's Willow processor has provided a major peer-reviewed result in error correction below threshold (Acharya et al., 2025), Microsoft's Majorana 1 has advanced a topological route based on a semiconductor-superconductor platform and a claimed topological core (Nayak, 2025), and Intel's silicon spin-qubit programme has demonstrated wafer-scale CMOS-compatible fabrication and cryogenic probing across 300 mm wafers (Neyens et al., 2024).

The central claim of this paper holds that no single architecture has yet solved the entire problem of future computation. Instead, different physical regimes appear to solve different bottlenecks. Biology broadens the conceptual horizon by showing that quantum-relevant organisation may occur in warm and noisy matter (Cao et al., 2020). Photonics offers a route beyond resistive interconnect loss (Bonaccorso et al., 2010). Superconducting codes currently provide the strongest experimental evidence for scalable logical suppression of error (Acharya et al., 2025). Topological qubits offer an intellectually powerful but still contested route to hardware-level protection (Nayak, 2025; Castelvecchi, 2025). Silicon spin qubits offer perhaps the strongest manufacturing logic because they align with the industrial capabilities of advanced semiconductor fabrication (Neyens et al., 2024).

Introduction

For much of the modern era, computing advanced through silicon scaling, improved lithography, and ever smaller transistors. That trajectory now encounters two different physical constraints. In classical and AI systems, one major bottleneck lies in interconnects, where resistance, power dissipation, and data movement increasingly limit performance. In quantum systems, the corresponding bottleneck lies in decoherence, noise, and the overhead required to convert imperfect physical qubits into reliable logical ones. As a result, the future

of computation depends less on a single successor to silicon and more on how materials, light, and device physics can be reorganised to address transport, error, and manufacturability together (Bonaccorso et al., 2010; Acharya et al., 2025).

At the same time, recent work in quantum biology has reopened a deeper question. If engineers must go to great lengths to preserve quantum coherence in artificial devices, what should be made of biological systems that may already exploit non-trivial quantum effects under ambient conditions? Reviews of the field now treat quantum biology as a legitimate scientific area rather than a fringe curiosity, while also distinguishing better-established biological cases from more speculative extensions (Cao et al., 2020). This does not mean that life should simply be equated with a quantum computer. It does mean that biology now belongs in the intellectual background of future computation, because it challenges older assumptions about what warm, noisy matter can or cannot do (Cao et al., 2020).

This dissertation therefore proceeds in four stages. First, it examines quantum biology as a prehistory of future computation. Second, it analyses graphene photonics as a response to the transport crisis in classical and AI hardware. Third, it evaluates Google's Willow result as the clearest recent systems-level advance in quantum error correction. Fourth, it compares Microsoft's Majorana programme and Intel's silicon spin-qubit work as alternative routes to scalable quantum devices. The aim is not to reduce these domains to one theory, but to show that they belong to one broader transformation in the understanding of matter, light, and information (Bonaccorso et al., 2010; Cao et al., 2020; Acharya et al., 2025; Neyens et al., 2024).

Quantum Biology as a Prehistory of Future Computation

Quantum biology concerns the possibility that living systems may preserve or exploit quantum effects in functionally relevant ways. Cao and colleagues identify several important areas in which this question has matured, including excitonic energy transfer in photosynthesis, radical-pair mechanisms in magnetoreception, and quantum tunnelling in enzymatic processes (Cao et al., 2020). Their review also shows that the field has become more disciplined, moving away from overstatement and toward a clearer separation between experimentally grounded cases and speculative interpretation (Cao et al., 2020).

This matters for computation because it shifts the conceptual baseline. Conventional engineering often treats quantum coherence as an exceptional achievement requiring isolation, cryogenics, and careful shielding. Quantum biology raises the possibility that certain forms of biological organisation can preserve useful quantum features through structural coupling and dynamical order rather than sterility alone (Cao et al., 2020). In that sense, biology becomes relevant not merely as a domain for scientific explanation but as a possible source of design intuition for future hardware.

The wider historical backdrop often returns to Schrödinger's *What is Life?*, which helped establish the question of whether life demands a subtler physical description than ordinary thermodynamic disorder would suggest. Modern retrospectives treat this work as an important part of the genealogy of quantum biology, even though the contemporary field has moved far beyond Schrödinger's original framework (Cao et al., 2020). Thus, the biological strand enters the present dissertation not as metaphor but as an academic precursor to present debates about coherence, information, and material organisation.

Biological Coherence and the Computational Capacity of Life

A more ambitious extension of the biological argument appears in Philip Kurian's 2025 work on the computational capacity of life. Kurian frames life in relation to physical limits on information processing and argues that one can meaningfully ask whether carbon-based life, taken collectively across Earth's history, has performed computation on a scale of cosmological significance (Kurian, 2025). This argument belongs in the dissertation because it directly joins biological systems to the language of computational capacity, speed limits, and large-scale informational organisation.

The importance of this work lies less in any single maximal claim than in its expansion of the question itself. If biological systems exhibit collective optical or quantum behaviour in protein-rich networks, then computation no longer appears as something that begins only with nervous systems or human-built machines. Biology starts to look like a distributed physical information processor, even if the exact meaning and extent of that phrase remain actively debated (Kurian, 2025). That shift places biological matter into the same conceptual field as photonic and quantum hardware, all of which concern how physical structure constrains information flow and transformation.

A serious academic treatment should still preserve a distinction between established field-level quantum biology and more ambitious claims about the aggregate computational significance of all life. The former enjoys a stronger consensus basis in the literature (Cao et al., 2020). The latter occupies a frontier position. Yet frontier status does not justify omission. It justifies careful placement. In the present dissertation, Kurian's work therefore functions as a speculative extension grounded in an identifiable academic lineage rather than as a settled conclusion.

Microtubules, Consciousness, and Speculative Extensions

The most controversial extension of quantum biology concerns consciousness and microtubules. This tradition, often associated with Penrose and Hameroff, proposes that intracellular microtubular structures may play a role in consciousness through quantum processes not captured by ordinary accounts of classical neural firing. Although such claims were long treated as highly marginal, the debate remains active enough that it cannot be dismissed out of hand within a dissertation on future computation. In particular, recent work

in *Neuroscience of Consciousness* has argued that a quantum microtubule substrate can address the binding and epiphenomenalism problems and has presented itself as experimentally supported (Wiest, 2025).

This does not establish consensus. It does, however, establish academic presence. The significance for this dissertation lies in the way such arguments blur the boundary between biological information processing and engineered quantum computation. If structured intracellular matter can no longer be ruled out as informationally richer than once assumed, then biology and quantum hardware begin to mirror one another conceptually. Both become questions about how physical systems can protect, organise, and exploit subtle states of information. The claim remains more speculative than the better-established areas of quantum biology, but it clearly links to the same underlying theme.

For that reason, the microtubule-consciousness literature belongs here in a carefully bounded form. It should neither dominate the dissertation nor be omitted. Its proper role lies in marking the outer edge of the biological argument, where the study of life's internal informational organisation begins to overlap with questions traditionally reserved for the philosophy of mind and quantum foundations.

The Return of the Interconnect Problem

If quantum biology broadens the conceptual frame, graphene photonics addresses a pressing engineering constraint. In advanced classical and AI hardware, one of the main performance bottlenecks now lies not only in logic devices but in the movement of data between them. As copper wires become thinner and traffic volumes increase, resistive losses, heat, and delay become increasingly costly. This makes interconnects a central problem rather than an afterthought of system design (Bonaccorso et al., 2010).

Photonics responds to this problem by replacing some electronic transport with optical transport. Light can carry information with lower resistive loss than copper, and it can support wavelength multiplexing, allowing multiple data streams to share the same physical pathway. That makes photonics attractive not only for long-distance communication but also for on-chip and chip-to-chip interconnects in data-intensive environments (Bonaccorso et al., 2010). In this sense, the future of classical computing depends not just on faster switching but on rethinking the physical medium of information flow.

Graphene as a Photonic Material

Graphene has attracted attention in this context because it combines high carrier mobility, optical transparency, mechanical flexibility, and broadband optical response in a one-atom-thick carbon layer. Bonaccorso and colleagues argued that graphene's true potential may lie in photonics and optoelectronics, precisely because its optical and electronic properties can be exploited even without a conventional semiconductor bandgap (Bonaccorso et al., 2010).

That argument remains important because it shifts graphene from the status of a fashionable material to that of a strategically useful photonic platform.

The crucial property here is tunability. Graphene interacts with light over a wide spectral range, and its absorption and related optical behaviour can be altered electrically. This makes it an attractive candidate for modulators, switches, photodetectors, and other compact integrated optical devices (Bonaccorso et al., 2010). Compared with conventional silicon photonics alone, graphene offers a broader and potentially more flexible materials basis for optical control.

This is one reason the graphene strand belongs beside the biological strand. In both cases, the question concerns what highly organised matter can do with information. Biology asks this in the context of living systems. Graphene photonics asks it in the context of engineered chips. The domains are not equivalent, but they both concern the informational potential latent in structure.

Optical Modulation and Computation with Light

Once light can be modulated on-chip, computation itself becomes possible in optical form. Interferometric optical systems can encode input values, weights, and phase shifts into light and can implement linear algebra operations central to AI workloads. This matters because matrix multiplication dominates large portions of modern machine learning, and photonic systems can in principle perform such operations with high throughput and low transport cost. Graphene strengthens this possibility by improving the modulation layer and shrinking the control footprint (Bonaccorso et al., 2010).

The significance is architectural. Photonics no longer appears only as a replacement for some wires. It begins to enter the compute plane itself. Once that happens, the distinction between communication hardware and compute hardware becomes less rigid. This is one reason graphene photonics belongs in a dissertation on the future of computation rather than in a narrower materials survey. It addresses both movement and transformation of information.

The Manufacturing Problem in Graphene Photonics

The weakness of graphene has long been manufacturability rather than conceptual promise. Materials that look extraordinary in laboratory-scale devices often encounter difficulties at wafer scale because defects, impurities, and transfer losses degrade performance. Graphene remains vulnerable to exactly these constraints, which is why its commercial future depends on fabrication infrastructure and yield rather than on demonstrations alone (Bonaccorso et al., 2010).

This is why Black Semiconductor's FabONE project is noteworthy. The company states that it is constructing a scalable 300 mm pilot line, aiming for operation after construction, pilot production thereafter, and later transition to volume production. It also explicitly frames its

work as graphene-based chip connectivity for the next generation of computing (Black Semiconductor, 2025). These claims do not prove that graphene photonics will dominate future hardware, but they do show that the field has entered a serious industrial phase.

Quantum Computing and the Problem of Error

If photonics addresses the transport limits of classical systems, quantum hardware addresses the problem of computational regimes unavailable to classical machines. Yet quantum computing remains dominated by the challenge of error. Useful quantum computation requires logical qubits whose effective error rate falls as protection increases. Without that property, larger systems simply create more opportunities for failure (Acharya et al., 2025).

This makes error correction more important than raw qubit counts alone. A large but noisy array does not by itself produce a useful machine. What matters is whether the architecture can cross the threshold at which adding more corrected structure suppresses errors rather than amplifying them. For many years, that remained one of the central unresolved tests of the field.

Google Willow and Error Correction Below Threshold

Google's Willow result matters because it offers a strong peer-reviewed answer to that problem. The Nature paper reports two below-threshold surface-code memories on Willow processors, including a distance-7 code and a distance-5 code integrated with a real-time decoder. It further reports that the logical error rate of the larger memory is suppressed when the code distance increases, culminating in a 101-qubit distance-7 code with a reported error per cycle of $0.143\% \pm 0.003\%$ (Acharya et al., 2025).

The importance of this result lies in what it changes experimentally. It shows that error correction can begin to work in the intended direction at meaningful scale. That does not amount to a practical universal quantum computer, but it does move the field beyond a purely aspirational claim. For the purposes of this dissertation, Willow therefore functions as the clearest recent benchmark in systems-level quantum progress (Acharya et al., 2025).

Microsoft Majorana and the Topological Proposition

Microsoft's Majorana programme approaches the same problem from a different angle. Rather than relying mainly on large correction overheads layered atop fragile qubits, it seeks a qubit architecture with stronger intrinsic protection. Microsoft's own announcement describes Majorana 1 as being built on a "topoconductor", a semiconductor-superconductor materials platform intended to create a topological state suitable for topological qubits, and presents this as a route toward practical large-scale quantum systems (Nayak, 2025).

The attraction of such a strategy is obvious. If quantum information can be stored non-locally and thereby made less sensitive to local noise, the cost of error correction could in principle

fall substantially. A topological qubit would thus represent a more elegant and potentially more scalable physical starting point than many existing approaches (Nayak, 2025).

Yet the academic standing of this result requires caution. Nature reported scepticism from physicists about whether Microsoft had yet demonstrated the full evidentiary basis for its topological-qubit claim, and later reporting indicated that important questions remained about the protocol underlying the claim (Castelvecchi, 2025; Garisto, 2025). APS Physics likewise framed the announcement as one that prompted expert scrutiny rather than consensus acceptance (Ball, 2025). The Majorana result therefore belongs in this dissertation as a serious but contested milestone rather than as a closed breakthrough on the level of Willow's error-correction result.

Intel Silicon Spin Qubits and CMOS Compatibility

Intel's silicon spin-qubit approach offers a third route, based not on superconducting surface-code dominance or topological protection, but on the manufacturability of semiconductor-compatible qubits. In the spin-qubit model, quantum information is encoded in the spin state of a single electron confined within a silicon quantum dot. The attraction lies partly in qubit compactness and partly in process compatibility with modern semiconductor fabrication.

Neyens and colleagues reported a cryogenic 300 mm wafer-probing process that enabled high-volume data collection on industry-manufactured spin-qubit devices at 1.6 K. They also reported that this process supports optimisation of CMOS-compatible fabrication, high yield, and low process variation across full wafers (Neyens et al., 2024). Intel separately described its Tunnel Falls device as a 300 mm silicon spin-qubit research chip built with advanced industrial techniques, explicitly positioning it as a way to accelerate research toward scalable silicon quantum systems (Intel, 2023).

This matters because manufacturing often decides the long-term fate of a technology. A qubit architecture that can exploit mature semiconductor infrastructure enjoys an obvious strategic advantage if the remaining control and fidelity problems can be solved. Silicon spin qubits therefore represent perhaps the clearest bridge between quantum computing and the industrial legacy of classical chip fabrication.

Comparative Assessment

A comparison across these domains suggests that future computation will be plural rather than singular. Quantum biology broadens the theoretical background by showing that living systems may exploit quantum-relevant organisation under ambient conditions (Cao et al., 2020). Graphene photonics addresses the transport crisis in classical and AI systems by moving information into the optical domain and by exploiting graphene's tunable material properties (Bonaccorso et al., 2010). Google's Willow establishes the strongest recent peer-reviewed evidence for below-threshold quantum error correction in a superconducting

platform (Acharya et al., 2025). Microsoft's Majorana programme offers the boldest route to hardware-level protection, but one whose strongest claims remain under active scrutiny (Nayak, 2025; Castelvechi, 2025; Garisto, 2025). Intel's silicon spin qubits offer the most industrially legible manufacturing pathway if their system-level maturity can continue to improve (Neyens et al., 2024).

The most plausible medium-term future therefore involves coexistence. Photonics is likely to strengthen interconnects and selected linear-algebra workloads. Superconducting systems are likely to continue leading in demonstrated error-corrected performance. Topological systems may yet prove decisive if the experimental case becomes more secure. Silicon spin qubits may become especially important where density, repeatability, and fab compatibility matter most. Biology will remain the conceptual horizon against which all these developments continue to be re-read, because it presses the question of whether the deepest informational architectures were present in nature before they were engineered in the laboratory.

Conclusion

The future of computation now unfolds across multiple physical fronts at once. One front concerns how information can move without being trapped by resistance, heat, and bandwidth limits. Another concerns how quantum states can survive long enough to support useful logical operations. A third concerns whether biology already offers examples, precedents, or at least suggestive analogues of informational organisation that engineers are only now learning to stabilise and reproduce. Taken together, these fronts show that the next era of computation will not be explained adequately by transistor scaling alone (Bonaccorso et al., 2010; Cao et al., 2020; Acharya et al., 2025).

The broad conclusion of this dissertation is therefore methodological as much as technological. Future computation will likely belong not to one miraculous device or one triumphant material, but to whichever stack best combines physical insight, error management, transport efficiency, and industrial repeatability. In that broader account, quantum biology restores depth to the story, graphene photonics restores attention to movement and energy, Willow restores empirical confidence in logical suppression of error, Majorana restores the ambition of hardware-level protection, and silicon spin qubits restore the manufacturing logic that has shaped every successful computing platform to date.

References

- Acharya, R. et al., (2025) *Quantum error correction below the surface code threshold*. London, Springer Nature.
- Ball, P., (2025) *Experts Weigh in on Microsoft's Topological Qubit Claim*. College Park, American Physical Society.

Black Semiconductor, (2025) *Welcome to our new headquarters, FabONE*. Aachen, Black Semiconductor.

Bonaccorso, F., Sun, Z., Hasan, T., Ferrari, A. C., (2010) *Graphene photonics and optoelectronics*. London, Nature Publishing Group.

Cao, J., Cogdell, R. J., Coker, D. F., Duan, H. G., Hauer, J., Kleinekathöfer, U., Jansen, T. L. C., Mančal, T., Miller, R. J. D., Ogilvie, J. P., Prokhorenko, V. I., Renger, T., Tan, H. S., Tempelaar, R., Thorwart, M., and others, (2020) *Quantum biology revisited*. Washington, DC, American Association for the Advancement of Science.

Castelvecchi, D., (2025) *Microsoft claims quantum-computing breakthrough, but some physicists are sceptical*. London, Springer Nature.

Garisto, D., (2025) *Microsoft quantum-computing claim still lacks evidence: physicists are dubious*. London, Springer Nature.

Intel, (2023) *Intel's New Chip to Advance Silicon Spin Qubit Research*. Santa Clara, Intel Corporation.

Kurian, P., (2025) *Computational capacity of life in relation to the universe*. Washington, DC, American Association for the Advancement of Science.

Nayak, C., (2025) *Microsoft unveils Majorana 1, the world's first quantum processor powered by topological qubits*. Redmond, Microsoft.

Neyens, S. et al., (2024) *Probing single electrons across 300-mm spin qubit wafers*. London, Springer Nature.