

Lifecycle of a Fault-Tolerant Quantum Computer

The

Lifecycle of a Fault-Tolerant Quantum Computer

Author: Rowan Brad Quni-Gudzinis **Contact:** rowan.quni@outlook.com **ORCID:** 0009-0002-4317-5604 **ISNI:** 0000000526456062 **URL:** <https://papers.qnfo.org/papers/lifecycle-of-a-fault-tolerant-quantum-computer/> **Date:** 2025-12-20 **Version:** 1.0

The operational lifetime of a fault-tolerant quantum computer is fundamentally a race between the inevitable, continuous degradation of its physical hardware and the escalating resource cost of the classical computation required for the error correction to compensate. This dynamic process defines a finite window of viability for any given quantum device. The initial quality of the processor's qubits, defined by its system-environment coupling strength at manufacture (γ_0), does not determine whether this race will occur, but rather sets the starting line and influences the pace. Ultimately, every quantum computer is on a trajectory toward eventual obsolescence, either through physical failure when its error rate exceeds the fault-tolerance threshold, or through resource exhaustion as the cost of its own maintenance becomes computationally unsustainable.

1.0 Introduction

1.1 Decoherence as Primary

Constraint

The pursuit of scalable quantum computation is fundamentally a contest against the phenomenon of decoherence, the irreversible loss of quantum information to the surrounding environment. This process represents the primary and most formidable constraint on the development of functional, large-scale quantum processors, acting as a universal limit on the duration and complexity of any quantum algorithm (Schlosshauer, 2007). Unlike the discrete, correctable errors of classical computing, decoherence is a continuous process that erodes the very foundation of quantum advantage: the delicate phase relationships inherent in superposition and entanglement. The viability of any quantum architecture is therefore not measured by its peak theoretical speed, but by the number of coherent operations it can perform within the finite window of time before its quantum state collapses into a classical one. Understanding, characterizing, and ultimately mitigating decoherence is not merely one aspect of quantum engineering; it is the central problem that defines the field.

The historical evolution of quantum theory has progressively moved from idealized, closed systems to the more realistic framework of open quantum systems, which explicitly acknowledges the impossibility of perfect isolation (Breuer & Petruccione, 2002). In the early decades of quantum mechanics, theoretical models often treated systems as if they were perfectly shielded from the universe, an abstraction that allows for elegant solutions but fails to capture physical reality. However, as the ambition to build functional quantum devices emerged in the late 20th century, it became clear that no system can be truly closed; every physical qubit inevitably interacts with its surroundings through thermal, electromagnetic, and acoustic channels. This unavoidable interaction is the physical origin of decoherence. The modern understanding, therefore, is that all quantum systems are open systems, and their temporal evolution must be described by mathematical tools that account for the continuous leakage of information into the environment.

The physical mechanism of decoherence is the formation of entanglement between the quantum system and the myriad, unobserved degrees of freedom of its environment (Schlosshauer, 2007). When a qubit in a superposition state interacts with its surroundings, information about its state is imprinted upon the environmental particles, such as photons, phonons, or nearby atomic spins. This process effectively creates a more complex, entangled state involving both the system and the environment. Because it is practically impossible to track and measure the state of every particle in the environment, the information that leaks out is considered lost. From the perspective of an observer who only has access to the qubit, this loss of information manifests as a decay of the off-diagonal elements of the system's reduced density matrix, a process that smoothly transforms a pure quantum superposition into a probabilistic classical mixture.

The computational model central to this analysis quantifies the effect of decoherence through the concept of a physical error rate, denoted as ϵ_p . This parameter represents the probability that a single quantum gate operation will fail due to an interaction with the environment. The model posits a direct relationship between the coherence time of a qubit and this error rate, formalizing the intuition that a shorter coherence window allows for fewer reliable operations. In this framework, decoherence is not an abstract concept but a measurable quantity that directly impacts computational fidelity. The entire structure of quantum error correction is built upon the necessity of combating this physical error rate, making its accurate characterization the first and most critical step in assessing the viability of any quantum processor.

A persistent counter-argument, often rooted in a purely theoretical perspective, suggests that decoherence could be entirely eliminated through the engineering of a perfectly isolated system. In this idealized view, if a quantum processor could be placed in a perfect vacuum, shielded from all electromagnetic fields, and cooled to absolute zero, it would form a closed system and maintain its coherence indefinitely. Proponents of this line of reasoning might argue that decoherence is therefore a technological problem of insufficient shielding, rather than a fundamental constraint. This perspective frames the challenge as one of achieving progressively better isolation, with the ultimate goal of removing the environment from the equation entirely (Breuer & Petruccione, 2002).

While the pursuit of better isolation is a critical engineering goal, the synthesis of theory and experiment confirms that perfect isolation is a physical impossibility. The laws of thermodynamics and quantum field theory ensure that no system can be completely decoupled from its surroundings; even in the deepest vacuum of space, a system is still bathed in the cosmic microwave background and subject to vacuum fluctuations. Furthermore, a useful quantum computer must have control and readout lines to receive instructions and report results, and these very lines act as unavoidable conduits for noise to enter the system. Therefore, the modern paradigm accepts that decoherence can never be eliminated, only managed (Schlosshauer, 2007). The central engineering challenge is not the futile quest for perfect isolation, but the practical task of making the coherence time as long as possible.

Accepting decoherence as an unavoidable and primary constraint immediately raises the critical follow-on question: what physical parameters govern the *rate* at which this information loss occurs? The answer is not a single, universal constant but is instead dependent on a nuanced interplay between the properties of the quantum system and its specific environment. The naive assumption that temperature is the sole determinant of this rate is a profound oversimplification that has historically led to incorrect conclusions about the feasibility of certain quantum technologies. A more physically accurate model must account for the specific nature of the interaction between the system and its environment, leading to a more complex but predictive framework.

1.2 Coupling

Strength and Temperature Interplay

The rate of decoherence in a quantum system is not determined by environmental temperature alone, but is more accurately described as being governed by the product of the system-environment coupling strength, denoted by the dimensionless parameter gamma (γ), and the temperature (T). This corrected physical model is the cornerstone of modern open quantum systems theory and the central premise of this revised analysis (Caldeira & Leggett, 1983). The coupling strength, γ , is a phenomenological constant that quantifies the intrinsic propensity of a qubit to interact with its surroundings, a property determined by its physical design, material composition, and local environment. Temperature, in contrast, characterizes the average energy of the environmental modes that the qubit can interact with. Both factors are equally critical, and a failure to account for their interplay leads to a fundamentally flawed understanding of quantum viability.

The historical focus on temperature as the primary enemy of quantum coherence led to the widespread, but ultimately incorrect, conclusion that room-temperature quantum computation was a physical impossibility for any and all systems. This “temperature-only” fallacy arises from a simple thermodynamic intuition: higher temperatures mean more environmental noise, which should overwhelm any quantum effect. While this intuition is not entirely wrong, it is incomplete. It overlooks the fact that the noise must be able to affect the qubit. The evolution of the field has been a gradual shift from this simplistic view to a more nuanced understanding, recognizing that a qubit that is very weakly coupled to its environment can remain coherent even in a high-temperature setting, much like a conversation in a soundproof room can remain clear even next to a noisy street.

The physical mechanism underlying the interplay between coupling and temperature is rooted in the dynamics of energy exchange. The environmental temperature, T , determines the thermal occupation number of environmental modes (such as photons or phonons) at a given frequency; a higher temperature means a greater number of thermal excitations are available to interact with the qubit. The coupling strength, γ , determines the probability that such an interaction will actually occur during a given period. The overall decoherence rate is therefore proportional to the product of the number of available noise sources (a function of T) and the strength of the interaction with those sources (γ). In the high-temperature limit, this relationship is approximately linear, with the decoherence rate scaling as γT , as first formalized in the canonical Caldeira-Leggett model (Caldeira & Leggett, 1983).

The numerical simulation at the heart of this work is explicitly designed to investigate this interplay, using a fixed cryogenic temperature of 0.02 Kelvin and varying the initial coupling strength, γ_0 , as the primary axis of analysis across seven different models. This approach isolates the effect of coupling on device viability. The most compelling real-world evidence for this model comes from two starkly different systems. The nitrogen-vacancy center in diamond exhibits millisecond-scale coherence times at room temperature, a feat possible only because the rigid diamond lattice provides exceptional isolation, resulting in a very weak γ (Doherty et al., 2013). Similarly, the nuclear spins used in clinical magnetic resonance imaging (MRI) maintain coherence for seconds inside the 310 Kelvin environment of the human body, an even more extreme example of a system with almost negligible coupling to its thermal surroundings (Pooley, 2005).

A counter-argument might posit that for the most promising scalable platforms, such as superconducting qubits, the coupling strength is intrinsically strong, and therefore extreme cooling is the only viable strategy to suppress decoherence. From this perspective, while the γT relationship is theoretically correct, the practical reality of engineering these devices makes T the only meaningful variable that can be controlled. Proponents of this view would argue that for any practical device, γ is a fixed property of the manufactured hardware, leaving temperature as the sole tunable parameter for improving coherence. This line of reasoning suggests that focusing on materials science to reduce γ is a less fruitful path than investing in more powerful cryogenic infrastructure.

The synthesis of theory and experimental practice reveals that both temperature and coupling are critical, tunable parameters. While it is true that achieving ultra-low temperatures is a non-negotiable requirement for many leading qubit modalities, it is not sufficient for achieving long coherence. Even at millikelvin temperatures, a

superconducting qubit that is strongly coupled to its environment—for example, through a high density of two-level system defects in its materials—will decohere rapidly (Müller et al., 2019). The decades-long effort to improve superconducting qubit coherence has been a two-front war: building better refrigerators to lower T , and simultaneously pioneering new fabrication techniques and material treatments to reduce γ . The ultimate viability of a quantum processor depends on the successful optimization of both.

This understanding of decoherence as a function of both initial coupling and temperature provides a static snapshot of a device's quality at the moment of its creation. However, a critical and often overlooked aspect of system viability is that these parameters, particularly the coupling strength, are not static over the operational lifetime of the device. Just as classical hardware ages and degrades, the physical substrate of a quantum processor is subject to changes that can alter its interaction with the environment, introducing a crucial temporal dimension to the problem of quantum viability.

1.3 Material Degradation

as Temporal Axis

The effective system-environment coupling strength, γ , of a quantum processor is not a fixed constant but rather a dynamic parameter that degrades—that is, increases—over its operational lifetime due to the gradual accumulation of material defects. This temporal degradation represents a fundamental lifecycle constraint on quantum hardware, transforming the challenge of quantum computing from a static design problem into a dynamic race against the inevitable aging of the device (Müller et al., 2019). The initial quality of a qubit, defined by its coupling strength at time zero (γ_0), determines its starting performance, but the rate of degradation determines its useful operational lifespan. Any comprehensive model of quantum viability must therefore account for the temporal evolution of $\gamma(t)$ as a primary axis of analysis.

The concept of device aging is ubiquitous in classical electronics, where phenomena like electromigration and oxide breakdown limit the lifespan of integrated circuits. While quantum processors lack the high currents and moving parts of many classical systems, they are subject to a more subtle and insidious set of aging mechanisms. Operating in a constant state of bombardment from environmental radiation and subject to the stresses of repeated thermal cycling between cryogenic and room temperatures, the atomic-scale structure of the device is not static (Vepsäläinen et al., 2020). Over months and years of operation, these accumulated insults manifest as changes in the material properties that govern the qubit's interaction with its environment, leading to a slow but inexorable decline in performance.

The physical mechanism behind this degradation can be understood as the creation of new decoherence channels. For solid-state qubits, a primary source of coupling is the presence of microscopic two-level system (TLS) defects in the amorphous materials of the device (Müller et al., 2019). It is plausible that events such as the impact of a high-energy cosmic ray or the localized stress from thermal contraction can cause atomic rearrangements in the substrate, creating new TLS defects that were not present at the time of manufacture. Each new defect represents an additional pathway through which the qubit can lose energy to its environment, effectively increasing the overall coupling strength γ . This process, integrated over time, leads to a measurable increase in the decoherence rate, even if the operating temperature remains perfectly stable.

The computational model employed in this study formalizes this process using a first-order degradation model, where the coupling strength is assumed to increase exponentially over time according to the formula $\gamma(t) = \gamma_0 * \exp(k_{degrade} * t)$. The degradation constant, $k_{degrade}$, is set to $5e-8$ per second, a value chosen to represent a characteristic lifetime on the order of several months, consistent with anecdotal observations of device performance drift in long-term experiments. This exponential model captures the essential behavior of an

accelerating degradation process, where the accumulation of defects makes the system progressively more susceptible to further damage. This time-dependent $\gamma(t)$ serves as the core dynamic input for the subsequent analysis of the system's lifecycle.

A potential counter-argument is that a quantum device, being a solid-state system operated in a highly controlled, cryogenic vacuum environment, should be almost perfectly stable over time. Unlike a classical computer with fans and spinning hard drives, a quantum processor has no moving parts and is shielded from many common sources of wear and tear. From this perspective, one might assume that once a device is characterized, its parameters, including γ , should remain fixed indefinitely. This view would suggest that focusing on temporal degradation is a secondary concern compared to the primary challenge of improving the initial manufacturing quality of the device.

This synthesis of evidence, however, demonstrates that the quantum realm is far from static. At the atomic scale, the device is a dynamic environment. The unavoidable flux of environmental radiation, including cosmic-ray muons that can easily penetrate standard laboratory shielding, provides a constant source of high-energy impacts that can and do alter the material substrate (Vepsäläinen et al., 2020). Furthermore, the very act of operating the qubit with strong microwave control pulses can induce material changes over billions of cycles. Therefore, the assumption of a static γ is a non-physical idealization. The properties of the device inevitably evolve over time, and a robust system architecture must be designed to anticipate and accommodate this degradation.

The direct and unavoidable consequence of a temporally increasing coupling strength, $\gamma(t)$, is a corresponding increase in the physical error rate of the quantum processor. As the barrier isolating the qubits from their environment weakens, the probability of an error occurring during any given gate operation rises. This escalating error rate poses a direct threat to the integrity of any long computation and cannot be ignored. To have any hope of performing a useful calculation on such a degrading physical substrate, an active and adaptive defense mechanism is required.

1.4 Quantum Error

Correction as Active Defense

Quantum error correction (QEC) represents the primary and only known scalable strategy to actively combat the effects of decoherence and other physical errors, thereby enabling reliable computation on inherently noisy quantum hardware. It functions as an active defense mechanism that continuously monitors for and reverses the errors caused by the environment, including those arising from the temporal degradation of the device (Google Quantum AI, 2023). Rather than attempting the impossible task of building a perfect physical qubit, QEC provides a systematic architectural solution to bridge the gap between the performance of faulty physical components and the requirements of a fault-tolerant quantum algorithm. The implementation of QEC transforms the challenge from building a perfect device to building one that is “good enough” for the correction code to be effective.

The concept of QEC has been a cornerstone of quantum information theory since the mid-1990s, but it is only in the current “Noisy Intermediate-Scale Quantum” (NISQ) era that experimental systems have reached the scale and fidelity required to begin implementing these codes in practice (Preskill, 2018). The historical progression has moved from theoretical possibility to active, real-time experimental demonstration. This shift marks a critical maturation of the field, moving beyond the characterization of single, isolated qubits to the systems-level challenge of making a collection of interacting, noisy qubits behave as a single, reliable computational unit. The development of effective QEC is now the central focus on the critical path toward scalable, fault-tolerant quantum computation.

The core mechanism of QEC involves encoding the information of a single “logical” qubit redundantly across a larger number of “physical” qubits. This redundancy allows the system to detect errors without directly measuring—and thus destroying—the delicate quantum state of the logical qubit (Google Quantum AI, 2023). Specialized

“ancilla” qubits are used to periodically measure collective properties of the physical data qubits, a process which reveals an “error syndrome.” This syndrome indicates whether an error has occurred and, crucially, what type of error it was (e.g., a bit-flip or a phase-flip). This classical syndrome information is then processed by a classical computer, which determines the appropriate corrective operation to apply to the physical qubits to restore the original logical state.

The computational model in this analysis adopts the surface code, which is the leading QEC protocol for many of the most promising solid-state quantum computing platforms. The strength of the protection offered by the surface code is determined by its “code distance,” an odd integer d . A higher code distance provides greater protection against errors but requires more physical qubits. The model captures this relationship by dynamically calculating the required code distance $d(t)$ needed to suppress the evolving physical error rate $\epsilon_p(t)$ down to a constant, target logical error rate ϵ_L . This adaptive code distance represents the active response of the QEC system to the degradation of the underlying hardware.

A persistent counter-argument in the field is that it is more efficient to invest all available resources into building better physical qubits rather than relying on the vast overhead of QEC. Proponents of this “purity” approach argue that QEC is a brute-force solution that consumes an enormous number of qubits (d^2 per logical qubit) that could otherwise be used for computation. They might point to systems with intrinsically very low error rates, such as certain silicon spin qubits or trapped ions, and argue that perfecting these platforms is a more direct path to fault tolerance than compensating for the flaws of noisier platforms like superconducting qubits (Burkard et al., 2021).

The current scientific consensus, supported by both theory and landmark experiments, is that perfect physical qubits are a physical impossibility. While building the best possible physical qubits is essential, no single qubit will ever be perfect enough to run a large-scale quantum algorithm without error correction. QEC is the only known solution that is scalable—that is, the logical error rate can be suppressed to arbitrarily low levels simply by increasing the code distance. Recent experiments have successfully demonstrated that QEC can extend the lifetime of a logical qubit beyond that of its constituent physical parts, proving that the overhead is not merely wasted but provides a tangible benefit (Sivak et al., 2023). The two approaches are not in opposition; they are complementary. Better physical qubits make the demands on QEC less extreme, but they do not eliminate the need for it.

The protection afforded by quantum error correction, however, is not a free lunch. The act of encoding information across more qubits, performing more frequent syndrome measurements, and processing more complex syndrome data carries a significant and continuously escalating resource cost. This cost is not borne by the quantum processor itself, but by the classical computational systems that are required to control it, creating a critical link between the health of the quantum device and the demands placed on its classical support infrastructure.

1.5 Escalating Resource Cost of

QEC

The computational resource cost of implementing quantum error correction scales polynomially with the required code distance, which must itself increase to counteract the rising physical error rate of a degrading quantum processor. This creates a direct and unavoidable link between the physical health of the quantum device and the computational burden placed on its classical control system (Google Quantum AI, 2023). As the quantum hardware ages and becomes noisier, the classical system must work exponentially harder to maintain a stable level of performance. This escalating resource cost is a fundamental constraint on the operational lifetime and economic viability of a fault-tolerant quantum computer.

In the early, theoretical stages of quantum computing, the resource cost of QEC was often considered an abstract accounting of the number of physical qubits required. However, as the field has moved towards experimental implementation, the focus has broadened to include the immense classical computational challenge that QEC represents (Pauka et al., 2021). The task of decoding error syndromes and calculating corrective operations must be performed in real-time, within a fraction of the qubit coherence time. This has created a new sub-field of research focused on developing high-speed, specialized classical hardware and decoding algorithms capable of keeping pace with the quantum device, a challenge that grows more acute as the device degrades.

The mechanism behind this escalating cost has two primary components. First, the number of physical qubits required to encode a single logical qubit in the surface code scales quadratically with the code distance, as d^2 . Increasing the code distance from $d=11$ to $d=13$ does not just add a few qubits; it requires increasing the physical qubit count from 121 to 169. Second, the classical computation required to decode the error syndromes also becomes more complex with increasing d (Sivak et al., 2023). The decoder must process a larger amount of syndrome data and search for the most likely error chain over a larger graph, a task whose complexity scales polynomially with the number of qubits.

The computational model in this study quantifies this overhead by calculating the resource cost, R_{cost} , as being directly proportional to the square of the practical code distance, $d_{\text{practical}}^2$. This serves as a robust proxy for the combined cost of physical qubit overhead and classical processing demands. The simulation logs clearly demonstrate this escalation: for the “Workhorse Device” (MODEL_03), the resource cost begins at 121 arbitrary units (for $d=11$), increases to 169 units after 187.5 days ($d=13$), and further increases to 225 units after 395.8 days ($d=15$). This represents an 86% increase in the computational cost of operation over a 14-month period.

A common counter-argument is that the relentless progress of classical computing, often colloquially described by Moore’s Law, will ensure that classical control systems can easily keep pace with the escalating demands of QEC. From this perspective, any increase in the R_{cost} of QEC will be trivially absorbed by the next generation of faster, more efficient classical processors (CPUs, FPGAs, or ASICs). This view suggests that the classical resource cost is not a significant long-term bottleneck, as classical technology will always advance faster than quantum hardware degrades.

This argument overlooks the unique and demanding nature of the QEC decoding problem. The challenge is not one of raw throughput, but of extremely low latency. The entire cycle of syndrome measurement, communication to the decoder, classical computation, and communication of the correction back to the QPU must be completed in a timescale much shorter than the qubit’s coherence time—typically on the order of microseconds or even nanoseconds (Sivak et al., 2023). This is not a task that can be offloaded to a distant supercomputer; it requires highly specialized, co-located hardware. While classical performance will improve, the polynomial scaling of the decoding problem combined with the stringent latency requirement means that the resource cost remains a very real and significant constraint, particularly as d grows large.

The escalating resource cost of QEC, driven by the degradation of the physical hardware, creates a dynamic tension that defines the operational lifecycle of the quantum computer. The system is viable only as long as it can successfully navigate this tension. This necessitates a formal definition of the boundaries of computational viability, establishing the precise conditions under which the system can be considered functional and the points at which it must be considered to have failed.

1.6 Defining Computational

Viability

The computational viability of a fault-tolerant quantum processor is not a binary state but a sustained condition bounded by two distinct failure thresholds: a physical limit defined by the efficacy of quantum error correction, and a resource limit defined by the capacity of the classical control system. A device is only viable as long as it operates within the envelope defined by these two constraints (Preskill, 2018). The operational lifetime of the machine is the duration of time it can maintain this state before the inevitable process of material degradation pushes it past one of these critical tipping points. This dual-constraint framework is essential for a realistic assessment of quantum computing feasibility.

The historical definition of quantum viability has evolved with the maturity of the field. In the early days of single-qubit experiments, viability was often simply the demonstration of coherence for a time longer than a single gate operation. As the field progressed into the NISQ era, the focus shifted to system-level metrics like quantum volume, which attempt to capture both the number and quality of qubits (Preskill, 2018). For the emerging era of fault tolerance, however, a more rigorous definition is required. Viability can no longer be a measure of average performance but must be a guarantee of sustained, reliable operation at a specified logical error rate, a guarantee that is contingent on both the quantum and classical components of the system.

The two failure conditions arise from different aspects of the QEC process. The first is the physical failure threshold. The threshold theorem of quantum error correction states that for a given code (like the surface code), there is a maximum physical error rate, ϵ_{th} , beyond which the code ceases to function (Google Quantum AI, 2023). If the physical error rate ϵ_p of the qubits exceeds this threshold, the QEC process itself introduces more errors than it corrects, leading to a catastrophic failure. The second is the resource failure threshold. The classical decoder engine is designed with a maximum computational capacity, R_{max} . If the required resource cost of QEC, $R_{cost}(t)$, which escalates as the device degrades, exceeds this capacity, the decoder cannot keep up with the quantum processor, leading to a backlog of uncorrected errors and an effective failure of the system.

The simulation suite at the core of this analysis explicitly models these failure conditions. The physical failure threshold is set at a standard value for the surface code, $\epsilon_{th} = 0.01$. The simulation for the “Sub-Threshold Device” (MODEL_07) demonstrates this failure mode: its initial physical error rate is calculated to be above 0.01, causing the simulation to terminate immediately at $t=0$ with a “Physical error rate exceeds threshold” message. While the simulation does not set an explicit R_{max} , the logs for the “Critical Threshold Device” (MODEL_06) illustrate the trajectory towards a resource failure, with the resource cost R_{cost} growing by 460% over the simulation period, a rate that would inevitably surpass any fixed classical capacity.

A potential counter-argument is that a system should be considered viable as long as it is capable of producing any computational result, regardless of the error rate. From this perspective, a high logical error rate does not constitute a “failure” but simply a lower quality of output. This view might be applicable to certain heuristic or variational algorithms in the NISQ era, where some degree of noise can be tolerated. However, this perspective is fundamentally incompatible with the requirements of fault-tolerant quantum computation.

For the large-scale algorithms that promise to solve intractable problems, such as Shor’s algorithm for factoring or quantum simulation for drug discovery, the final result must be correct with a very high degree of confidence. These algorithms are the entire motivation for building a fault-tolerant machine. A system that cannot guarantee a low logical error rate cannot correctly execute these algorithms and has therefore failed in its primary purpose (Sivak et al., 2023). The synthesis of theoretical requirements and experimental goals confirms that maintaining a logical error rate below a target value (e.g., $\epsilon_L = 1e - 15$ in the S3A model) is the non-negotiable definition of success. Breaching either the physical or resource threshold makes this guarantee impossible.

By establishing this rigorous, dual-constraint definition of viability, it becomes possible to frame the entire operational lifecycle of a quantum computer within a single, coherent narrative. This framework allows for a quantitative analysis and comparison of the useful operational lifetime of devices with different initial characteristics, moving beyond simple performance benchmarks to a more holistic assessment of long-term sustainability. This leads directly to the central thesis of this work: that quantum computation is an inevitable and continuous race against time.

1.7 Thesis: An

Inevitable Race Against Time

The operational lifetime of a fault-tolerant quantum computer is fundamentally a race between the inevitable, continuous degradation of its physical hardware and the escalating resource cost of the classical computation required for the error correction to compensate. This dynamic process defines a finite window of viability for any given quantum device. The initial quality of the processor's qubits does not determine whether this race will occur, but rather sets the starting line and influences the pace. Ultimately, every quantum computer is on a trajectory toward eventual obsolescence, either through physical failure or through the overwhelming cost of its own maintenance.

This perspective reframes the challenge of building a quantum computer from a static goal of achieving a certain performance metric to a dynamic, operational lifecycle problem. It shifts the focus from merely asking “how good is this device now?” to the more critical question of “for how long can this device maintain a state of computational viability?” This lifecycle perspective is essential for planning the long-term development of quantum infrastructure and for understanding the true economic and computational cost of sustaining a quantum advantage over time. It acknowledges that a quantum computer is not a permanent fixture but a high-performance machine with a finite operational lifespan.

The mechanism driving this race is a causal chain that links material science to computational complexity. The process begins with the slow accumulation of material defects, which causes the system-environment coupling strength, $\gamma(t)$, to increase over time. This increased coupling leads directly to a shorter coherence time and thus a higher physical error rate, $\epsilon_p(t)$. To maintain a constant, low logical error rate, the quantum error correction system must respond by increasing its protective power, which means increasing the code distance, $d(t)$. This, in turn, drives up the classical resource cost, $R_cost(t)$, required for real-time decoding. This feedback loop continues relentlessly until the system hits a wall: either $\epsilon_p(t)$ exceeds the physical fault-tolerance threshold, or $R_cost(t)$ exceeds the capacity of the classical control system.

The suite of seven simulations performed in this work was designed explicitly to model this entire lifecycle for devices of varying initial quality. The computational matrix spans a range of initial coupling strengths (γ_0) from the “Idealized Device” (MODEL01) to the “Sub-Threshold Device” (MODEL07). The numerical logs from these simulations provide a quantitative, time-resolved picture of this race. They show how a high-quality device with a low γ_0 experiences a slow, manageable increase in resource cost, while a low-quality device with a high γ_0 faces a rapid, unsustainable explosion in computational overhead, leading to a much shorter useful lifetime.

An optimistic counter-argument might propose that future breakthroughs in materials science will lead to the creation of perfectly stable quantum materials that do not degrade over time. In such a scenario, γ would remain constant, ϵ_p would be fixed, and the required QEC resources would never need to escalate. This would effectively halt the race before it begins, allowing for a quantum computer with an indefinite operational lifetime. This perspective places its faith in a future material science solution that would eliminate the problem of device aging entirely.

While the pursuit of more stable materials is a vital research direction, the fundamental principles of thermodynamics and the realities of operating complex devices in a radiation-filled universe suggest that perfect, indefinite stability is a physical impossibility. All known complex physical systems are subject to the relentless increase of entropy and the accumulation of defects over time. Therefore, the synthesis of current physical understanding and the simulation results leads to a more realistic conclusion: the initial quality of a device, its γ_0 , does not determine *if* it will fail, but rather *how long* it can sustain viable computation before it does. A lower γ_0 does not win the race; it simply allows the system to run for a longer time before exhaustion.

This thesis, which frames quantum viability as a finite and predictable lifecycle, provides the necessary context for the detailed analysis that follows. The subsequent sections of this paper will first detail the specific mathematical models and architectural components that underpin this simulation of a quantum computer's lifecycle. Following this, a thorough analysis of the numerical results will be presented, quantitatively comparing the operational lifetimes and resource cost trajectories for each class of device, from the idealized to the physically impossible.

2.0 Methodology

2.1 Open Quantum Systems

Framework

The computational model presented in this analysis is rigorously grounded in the standard theoretical framework of open quantum systems (Breuer & Petruccione, 2002). This approach is essential for ensuring the physical validity of the simulation, as it explicitly acknowledges that any real-world quantum processor is an open system, unavoidably interacting with its surrounding environment. By building upon this established foundation, the model guarantees that its core assumptions are consistent with the principles of quantum mechanics that govern decoherence and dissipation. The framework provides the mathematical language to connect abstract concepts like information loss to concrete physical parameters that can be measured in a laboratory and simulated numerically.

The historical development of quantum mechanics has seen a crucial evolution from the study of idealized, isolated “closed” systems to the more complex and realistic treatment of “open” systems. The Schrödinger equation, which describes the deterministic evolution of a closed system, is insufficient to capture the dynamics of a real qubit, which is subject to noise and thermal fluctuations from its environment. The modern framework of open quantum systems, developed throughout the latter half of the 20th century, provides the necessary theoretical tools, such as master equations and influence functionals, to describe the stochastic and dissipative evolution of a system in contact with a much larger, unobserved environment.

The central mechanism within the open quantum systems framework is the concept of a system coupled to a bath. The quantum processor is treated as the “system,” while the surrounding universe—the cryogenic apparatus, the control lines, the electromagnetic vacuum—is treated as the “bath.” The interaction between the system and the bath is what leads to decoherence (Breuer & Petruccione, 2002). The mathematical model formalizes this by relating the rate of decoherence to the properties of both the system and the bath. Specifically, the coherence time (T_{coh}) is shown to be inversely proportional to the strength of the system-bath interaction (the coupling, γ) and a function of the thermal energy of the bath (the temperature, T).

The Python code at the core of the S3A simulation directly implements this physical relationship. The inverse of the coherence time, which represents the decoherence rate, is calculated using the formula $T_{\text{coh}}^{-1} = (\gamma / (2 \pi k_B T)) / \hbar$. This line of code is a direct translation of the high-temperature limit approximation derived from the standard

open quantum systems framework. It explicitly links the output of the simulation to the foundational physics, ensuring that the model's behavior is not an arbitrary construction but a reflection of established quantum mechanical principles.

A potential counter-argument from a purely computational or engineering perspective might be that such a foundational physical model is unnecessarily complex. One could propose a more phenomenological, top-down model where the error rate is simply an empirical parameter measured from a device, without regard for the underlying physics of γ and T . Such a model might be simpler to implement and could still be used to project the resource costs of quantum error correction. This approach would prioritize empirical observation over theoretical grounding.

While a purely phenomenological model has its uses, grounding the simulation in the open quantum systems framework provides crucial predictive power and physical insight. By explicitly modeling γ and T , the simulation can make predictions about how a device's performance would change if the temperature were altered or if a new fabrication process reduced the coupling strength. This is essential for guiding future research and development. The synthesis of theory and application demonstrates that using the standard physical framework is not an unnecessary complication but a prerequisite for building a credible and predictive model of quantum viability (Breuer & Petruccione, 2002).

The standard open quantum systems framework typically treats the coupling strength γ as a static, time-independent parameter that characterizes a given device. However, to capture the long-term lifecycle of a quantum processor, this static assumption must be extended. The following section details the methodology for treating γ as a dynamic variable that evolves as a function of time, representing the physical degradation of the hardware.

2.2 First-Order

Degradation Model for Coupling

To account for the physical aging of the quantum processor, the temporal degradation of the system-environment coupling strength is modeled as a first-order exponential process. This mathematical choice posits that the rate of increase in coupling strength at any given time is proportional to the coupling strength that already exists. This approach provides a physically plausible and computationally tractable model for the accumulation of material defects and other sources of environmental noise over the operational lifetime of the device. It captures the essential behavior of a system that becomes progressively more susceptible to damage as it accumulates flaws.

First-order and exponential models are standard tools used across many scientific and engineering disciplines to describe processes of growth or decay where the rate of change is proportional to the current state. From population growth in biology to radioactive decay in physics, these models have proven to be remarkably effective at capturing the fundamental dynamics of a wide range of systems. In the context of material science, such a model can be interpreted as representing a process where the existence of defects (which contribute to γ) may act as nucleation sites or create local stress, making the formation of further defects more likely.

The mathematical mechanism of the model is straightforward. The system begins at time $t=0$ with an initial coupling strength γ_0 , which represents the pristine, as-manufactured quality of the device. The model then assumes that the rate of change of the coupling, $d\gamma/dt$, is equal to $k_{\text{degrade}} * \gamma(t)$, where k_{degrade} is a small, positive constant representing the intrinsic degradation rate. The solution to this simple differential equation is an exponential function: $\gamma(t) = \gamma_0 \exp(k_{\text{degrade}} t)$. This function describes a coupling strength that is initially stable but increases at a progressively faster rate over time.

The S3A simulation implements this exact formula in both its LaTeX theoretical framework and its Python execution code. A specific degradation constant of $k_{\text{degrade}} = 5\text{e-}8$ per second is chosen. This value is not arbitrary; it is selected to correspond to a “half-life” of approximately 231 days, meaning the coupling strength would double in that period. This represents a plausible timescale for significant performance drift in a complex, solid-state quantum device that is under continuous operation and environmental bombardment. The simulation logs for every model clearly show this exponential increase in $\gamma(t)$ at each time step.

A counter-argument could be made that this model is too simplistic and that the physical processes of degradation are likely to be far more complex. For example, one could argue that degradation might follow a power law, or that it might be characterized by discrete, stochastic jumps in γ corresponding to high-energy particle strikes, rather than a smooth exponential curve. A more complex model might also include terms for annealing or self-healing effects that could partially reverse the degradation process.

While it is true that the microscopic reality of degradation is undoubtedly more complex, the first-order exponential model serves as a robust and effective first approximation. It successfully captures the essential, qualitative behavior that is central to the thesis: that the device’s performance worsens over time, and that the rate of this worsening is not constant. For the purposes of a strategic, system-level analysis of a device’s lifecycle, this model provides the necessary dynamic input without getting bogged down in the yet-unknown microscopic details of specific degradation pathways. It represents a standard and effective choice for modeling aging processes in the absence of a more complete, experimentally verified theory.

The direct and most important consequence of this time-dependent coupling strength, $\gamma(t)$, is its impact on the physical error rate of the quantum processor. As the isolation of the qubits from their environment weakens over time, the probability of an error occurring during a gate operation must necessarily increase. The following section details the methodology used to formalize and calculate this crucial relationship.

2.3 Physical Error Rate

Formulation

The physical error rate, denoted as ϵ_p , is formulated as the dimensionless ratio of the time required to perform a quantum gate operation (t_{gate}) to the time available before the quantum state decoheres (T_{coh}). This definition provides a direct and physically intuitive link between the underlying physics of decoherence and the primary input parameter that governs the requirements of quantum error correction. It formalizes the concept that the reliability of a quantum operation is a competition between the speed of the control system and the speed of the environmental noise.

Historically, physical error rates in quantum computing have been characterized through a variety of experimental benchmarking protocols, such as randomized benchmarking, which yield an average error per gate. While these experimental values are crucial for characterizing a specific device, a predictive model requires a way to derive the error rate from more fundamental parameters. The formulation $\epsilon_p = t_{\text{gate}}/T_{\text{coh}}$ provides this necessary theoretical bridge. It allows the model to predict how the error rate will change in response to modifications in the system, such as using faster gates or improving the coherence time.

The mechanism captured by this formulation is straightforward. A quantum gate is an operation that takes a finite amount of time, t_{gate} , to execute. During this entire duration, the qubits are susceptible to decoherence, a process characterized by the timescale T_{coh} . The ratio of these two times can be interpreted as the probability that a decoherence event will occur at some point during the gate operation, thus corrupting the state and causing an error. As the device degrades, $\gamma(t)$ increases, which causes $T_{\text{coh}}(t)$ to decrease. Since t_{gate} is a fixed parameter of the control hardware, the ratio $\epsilon_p(t) = t_{\text{gate}}/T_{\text{coh}}(t)$ must therefore increase over time.

The S3A simulation code directly implements this calculation at every time step for each of the seven models. A fixed gate time of $t_{\text{gate}} = 4$ nanoseconds is used, which is a representative value for fast, state-of-the-art superconducting qubits. The coherence time, $T_{\text{coh}}(t)$, is continuously recalculated based on the current value of the degraded coupling strength, $\gamma(t)$. The resulting physical error rate, $\epsilon_p(t)$, is then logged and used as the input for the subsequent quantum error correction calculations. For example, in the “Workhorse Device” (MODEL_03), ϵ_p is shown to degrade from an initial value near 0.0000 to 0.0001 over the course of the 14-month simulation.

A valid counter-argument is that this formulation for the physical error rate is an oversimplification. In a real quantum processor, the total error rate is a composite of multiple independent error sources. In addition to decoherence during the gate (T_{coh} effects), there are also errors from imperfect gate calibration (gate infidelity), measurement errors, and state preparation errors. A more comprehensive model would treat ϵ_p as a sum of these various contributions, $\epsilon_p = \epsilon_{\text{decoherence}} + \epsilon_{\text{infidelity}} + \epsilon_{\text{measurement}} + \dots$

While acknowledging the existence of multiple error channels, the model’s focus on decoherence as the primary driver of the error rate is a methodologically sound simplification for a lifecycle analysis. Gate infidelities and measurement errors are typically addressed through intensive, periodic calibration routines and can be considered relatively static over the short term. Decoherence due to material degradation, however, represents a slow, continuous, and largely irreversible increase in the baseline error rate. By modeling ϵ_p as a direct function of $T_{\text{coh}}(t)$, the simulation captures the dominant dynamic component of the error budget over long timescales. It therefore serves as a robust proxy for the overall physical error rate’s temporal evolution.

The calculation of the time-dependent physical error rate, $\epsilon_p(t)$, is the final step in characterizing the degrading physical hardware. The subsequent challenge is to determine the necessary response from the quantum error correction system to counteract this rising tide of physical errors. The following section details the methodology used to model the error suppression capabilities of the surface code, the leading QEC protocol for this task.

2.4 Surface Code Error

Suppression Model

The methodology for modeling the effects of quantum error correction is based on the standard error suppression relationship for the surface code. This mathematical model describes how the logical error rate of a protected qubit (ϵ_L) is determined by the physical error rate of the underlying components (ϵ_p) and the “code distance” (d), a parameter that quantifies the strength of the error correction (Google Quantum AI, 2023). This formulation is critical as it provides the quantitative link between the quality of the physical hardware and the performance of the fault-tolerant logical qubit.

The surface code has emerged over the past two decades as the leading candidate for implementing fault-tolerant quantum computation, particularly for solid-state architectures like superconducting and semiconductor qubits. Its primary advantages are that it requires only nearest-neighbor interactions between physical qubits arranged on a 2D grid, and it possesses a remarkably high fault-tolerance threshold (Google Quantum AI, 2023). This means it can successfully suppress errors even with relatively noisy physical components. The historical development of the surface code from a theoretical concept to an experimentally implemented protocol provides the context for its selection as the basis for this model.

The error suppression mechanism of the surface code is based on topological principles. The logical information is encoded in a global, non-local property of the entire grid of physical qubits. Local errors, such as a single bit-flip on one physical qubit, can be detected by measuring local “stabilizer” operators. These measurements reveal an error syndrome that points to the location of the error without revealing the logical information itself. For an error to

corrupt the logical information, a chain of physical errors must occur that stretches all the way across the grid. The code distance, d , corresponds to the size of this grid. The probability of such a long error chain occurring by chance decreases exponentially as the code distance increases (Google Quantum AI, 2023).

This exponential suppression is captured by the standard scaling relationship, which is included in the S3A LaTeX framework: $\epsilon_L \approx C * (\epsilon_p / \epsilon_{th})^{((d+1)/2)}$. In this formula, C is a constant pre-factor, and ϵ_{th} is the fault-tolerance threshold, below which the code is effective. This equation shows that as long as ϵ_p is below ϵ_{th} , increasing the code distance d causes the logical error rate ϵ_L to decrease exponentially. The S3A simulation inverts this relationship to solve for the d required to achieve a target ϵ_L , given a certain ϵ_p .

A counter-argument from the perspective of theoretical computer science is that other families of quantum error-correcting codes exist, such as LDPC (Low-Density Parity-Check) codes, which offer potentially better scaling relationships. These codes might, in theory, achieve the same logical error rate with a smaller number of physical qubits (i.e., a lower overhead) compared to the surface code. Therefore, basing the entire model on the surface code might lead to an overly pessimistic estimation of the required resources.

While more efficient codes are an active and important area of research, the choice of the surface code for this model is methodologically justified by its practical advantages and widespread adoption in the experimental community. The high threshold ($\epsilon_{th} \approx 1\%$) and the requirement for only local connectivity make it far more practical to implement on today's 2D quantum processors than LDPC codes, which often require complex, long-range connections between qubits (Google Quantum AI, 2023). The synthesis of theory and practice shows that for the current and next generation of hardware, the surface code is the de facto standard, making it the most realistic and relevant choice for a model of near-term fault tolerance.

The surface code error suppression formula provides a static relationship between the physical and logical error rates for a fixed code distance. However, the core of this analysis is a dynamic simulation where the physical error rate is constantly changing. This requires a methodology for adaptively calculating the necessary code distance in real-time to maintain a constant level of logical performance, as detailed in the following section.

2.5 Required Code Distance

Calculation

To ensure the quantum computer maintains a constant, reliable level of performance over its operational lifetime, the methodology specifies that the quantum error correction code distance, d , must be dynamically increased to counteract the degradation of the physical error rate, $\epsilon_p(t)$. This is achieved by mathematically inverting the surface code error suppression formula to solve for the required code distance $d(t)$ as a function of the evolving physical error rate. This calculation represents the active, adaptive response of the QEC system to the changing health of the quantum hardware.

The concept of adaptive error correction is a standard feature of classical information and communication systems. For example, a cellular modem will automatically switch to a more robust but less efficient modulation scheme when the signal quality is poor. The methodology presented here applies the same principle to quantum computation. Instead of designing a system with a fixed, worst-case level of error correction, a more efficient approach is to design a system that can adapt its level of protection in response to the currently measured error rate, a concept that becomes crucial when considering a device whose error rate changes over its lifecycle.

The mathematical mechanism for this calculation is an algebraic rearrangement of the surface code scaling formula. Starting with $\epsilon_L \approx C * (\epsilon_p(t) / \epsilon_{th})^{((d+1)/2)}$, the goal is to solve for d . By taking the logarithm of both sides and rearranging the terms, one arrives at the expression for the required code distance:

$d(t) \approx 2 * [\ln(\epsilon_L/C) / \ln(\epsilon_p(t)/\epsilon_{th})] - 1$. This equation takes the current physical error rate $\epsilon_p(t)$ as an input and outputs the theoretical code distance $d(t)$ required to suppress that error rate down to the constant, target logical error rate ϵ_L .

The S3A Python simulation code implements this inverse calculation at every time step. It uses a target logical error rate of $\epsilon_L = 1e - 15$, a standard goal for fault-tolerant algorithms. The result of the formula, which is a real number, is then processed to find the practical code distance. Since the surface code distance must be an odd integer, the code calculates $d_practical = \text{math.ceil}(d / 2) * 2 + 1$, effectively rounding the theoretical value up to the next valid odd integer. The simulation logs for the “Workhorse Device” (MODEL_03) clearly show this adaptation in action, with $d_practical$ starting at 11 and being revised upwards to 13 and then 15 as $\epsilon_p(t)$ degrades.

A practical engineering counter-argument would be that a real quantum computing system cannot dynamically recompile its entire qubit layout to change the code distance on the fly. A system is likely to be designed and fabricated with a fixed, maximum code distance in mind. Therefore, modeling a continuously variable d might be physically unrealistic. The system would likely operate at a fixed d , and one would simply have to tolerate a logical error rate $\epsilon_L(t)$ that degrades over time.

This is a valid and important engineering consideration. However, the methodology of calculating a required $d(t)$ remains a powerful analytical tool. It can be interpreted in two ways. First, for a future, highly flexible architecture, it could represent a truly adaptive system. Second, and more immediately relevant, it serves as a requirement specification for a fixed-distance architecture. The calculation of $d(t)$ over the device’s target lifetime determines the *minimum* fixed code distance the device must be designed with to remain viable until its end-of-life. The calculation of a dynamic $d(t)$ is therefore the necessary first step in designing a static, but long-lived, fault-tolerant system.

The calculation of the required code distance, $d(t)$, is the final step in determining the necessary quantum response to hardware degradation. The direct and most significant consequence of a higher code distance is the increased demand on the system’s resources, particularly the classical computational power needed for decoding. The following section details the methodology used to model this escalating resource cost.

2.6 Resource Cost Function

The resource cost of implementing quantum error correction is modeled as a function that scales quadratically with the practical code distance, $d_practical$. This methodological choice is designed to capture the significant and rapidly growing overhead associated with deploying stronger levels of error protection. This cost function serves as a proxy for the combined demands on both the number of physical qubits and the classical computational power required for real-time decoding, providing the final output metric for evaluating the long-term sustainability of each simulated device.

The historical understanding of the cost of fault tolerance has always acknowledged a significant overhead, but the true scale of this cost has become more apparent as experimental implementations have been developed. Early theoretical work focused primarily on the qubit overhead, establishing the d^2 relationship for the surface code (Google Quantum AI, 2023). More recently, as real-time decoding has become an experimental reality, the immense challenge of the classical processing side of the problem has come into sharp focus. The resource cost is now understood to be a hybrid quantum-classical cost, dominated by the sheer number of physical components and the low-latency classical computation needed to manage them.

The quadratic scaling of the resource cost is rooted in the fundamental structure of the surface code. To build a logical qubit with code distance d , the number of physical data qubits required is d^2 , and the number of measurement qubits is $d^2 - 1$. The total number of physical qubits is therefore approximately $2d^2$ (Google Quantum AI, 2023). This is the quantum component of the cost. The classical component arises from the decoding

algorithm, which must process the results from $d^2 - 1$ measurements in each cycle. The complexity of the most efficient known decoders for the surface code, such as Minimum Weight Perfect Matching, scales polynomially with the number of measurements, often as a high-power polynomial like d^6 or d^8 . The choice of a d^2 scaling in the model is therefore a conservative lower bound on the true total resource cost.

The S3A Python simulation code implements this conservative cost function directly, calculating the resource cost R_{cost} as $d_{\text{practical}}^2$. This simple, powerful formula makes the consequences of degradation immediately apparent in the simulation logs. For the “Critical Threshold Device” (MODEL_06), the required code distance increases from $d=19$ to $d=41$ over the 14-month simulation. The resource cost, however, explodes from $19^2 = 361$ to $41^2 = 1681$ arbitrary units, a staggering 460% increase. This quadratic relationship turns a modest linear increase in the required protection into an unsustainable explosion in cost.

A counter-argument could be made that future advances in classical computing, particularly the development of specialized ASICs (Application-Specific Integrated Circuits) for QEC decoding, could significantly reduce the classical processing cost. If a highly efficient decoder could be designed with a complexity that scales more favorably than the qubit overhead, the overall resource cost might not be dominated by the d^2 term. This perspective suggests that innovation in classical hardware design could “flatten the curve” of the escalating resource cost.

While it is certain that specialized classical hardware will be a critical component of any fault-tolerant quantum computer, the physical qubit overhead of d^2 remains a fundamental and unavoidable property of the surface code itself. No amount of classical processing efficiency can reduce the number of physical qubits that must be fabricated, controlled, and measured in every cycle. Therefore, the quadratic scaling of the resource cost stands as a robust lower bound on the total system cost. The model’s use of $R_{\text{cost}} = d_{\text{practical}}^2$ correctly captures the dominant scaling factor that will drive the long-term economics and engineering of scalable quantum computers.

With the methodologies for modeling the physical hardware, its degradation, the required QEC response, and the resulting resource cost all established, the final step in the methodology is to define the specific initial conditions and parameters for the suite of simulations. This computational matrix is designed to sweep across a range of device qualities to provide a comprehensive analysis of the long-term viability of different classes of quantum processors.

2.7

Computational Matrix and Simulation Parameters

To systematically investigate the impact of initial hardware quality on the long-term viability and operational cost of a quantum processor, a computational matrix consisting of seven distinct models was constructed. Each model is defined by a unique initial system-environment coupling strength, γ_0 , representing a different class of device, from a future, idealized processor to one that is non-functional from the moment of fabrication. This matrix structure allows for a controlled, comparative analysis of the simulated lifecycles, isolating the effect of initial quality while holding all other system parameters constant.

The use of a computational matrix or a parameter sweep is a standard and essential methodology in computational science for exploring the behavior of a complex model. Instead of simulating a single, arbitrary data point, this approach maps out the model’s behavior across a wide range of input conditions. This is particularly important in a field like quantum computing, where the performance of devices can vary by orders of magnitude, from early-stage prototypes to state-of-the-art “hero” devices. The matrix is designed to capture this full spectrum of performance observed over the recent history and near-future projections of the field.

The seven models in the matrix are defined by their initial coupling strength, γ_0 , which serves as a proxy for their manufacturing quality. The values are chosen to span the range of interesting physical behaviors. MODEL_01 ($\gamma_0 = 1.0e - 8$) represents an “Idealized” device, better than current technology. MODEL_03 ($\gamma_0 = 2.0e - 7$) is the “Workhorse,” representing a typical, statistically average device. The matrix includes models for a “Hero” device (MODEL_02), a “Legacy” device (MODEL_04), and a “Noisy Prototype” (MODEL_05). Crucially, the matrix is designed to straddle the absolute limit of viability, with MODEL_06 ($\gamma_0 = 4.9e - 6$) at the “Critical Threshold” and MODEL_07 ($\gamma_0 = 6.0e - 6$) being “Sub-Threshold.”

To ensure a fair comparison, all other simulation parameters are held constant across the seven models. The operating temperature is fixed at $T = 0.02$ Kelvin, representing a standard dilution refrigerator environment. The gate time is fixed at $t_{\text{gate}} = 4$ nanoseconds, typical for fast superconducting qubits. The degradation constant is $k_{\text{degrade}} = 5e-8$ per second for all models. The QEC fault-tolerance threshold is $\epsilon_{th} = 0.01$. This strict control of variables ensures that any differences observed in the simulation outputs are solely attributable to the difference in the initial quality, γ_0 , of the device.

A potential counter-argument is that the specific choice of the seven γ_0 values is arbitrary and that the results might be different if other values were chosen. Furthermore, one could argue that a real-world analysis should also vary other parameters, such as the operating temperature or the gate speed, to explore a wider range of the design space. This would involve a much larger, multi-dimensional parameter sweep.

While a larger parameter sweep would certainly be informative, the chosen computational matrix is methodologically sound for the specific thesis of this paper. The seven γ_0 values were not chosen arbitrarily, but were carefully selected based on preliminary runs to ensure that they would illustrate the full range of qualitative behaviors, from long-term stability to immediate failure. By holding other parameters constant, the simulation provides a clear and unambiguous demonstration of the central role that initial hardware quality plays in determining the operational lifecycle of a quantum processor. It is a targeted experiment designed to test a specific hypothesis.

With the complete methodology for the simulation now established—from the foundational physical model to the specific initial conditions of the computational matrix—the stage is set for the presentation and analysis of the results. The data generated by executing this simulation provides the core quantitative evidence for the thesis of this work. The following sections will first describe the conceptual system architecture that this methodology models, before proceeding to a detailed analysis of the simulation logs for each of the seven device classes.

3.0 System Architecture

3.1 Quantum

Processing Unit and Initial Coupling

The architectural core of the simulated system is a Quantum Processing Unit (QPU), a solid-state chip containing an array of physical qubits. The single most critical parameter defining the QPU’s intrinsic quality and long-term potential is its mean system-environment coupling strength at the time of manufacture, denoted as γ_0 . This initial coupling is not an abstract variable but a direct physical consequence of the materials, design, and fabrication precision of the device (Burkard et al., 2021). A lower γ_0 signifies a higher-quality QPU with superior intrinsic isolation from environmental noise, which serves as the foundation for a longer and more efficient operational lifetime.

The history of quantum hardware development can be viewed as a continuous and painstaking effort to reduce the initial coupling strength of physical qubits. Early superconducting qubits from the late 1990s and early 2000s had coherence times in the nanosecond range, corresponding to a very high γ_0 . Over two decades, a series of

breakthroughs in qubit design (e.g., the transition from the Cooper-pair box to the transmon), material science (e.g., using higher-purity substrates and interfaces), and fabrication techniques (e.g., improved lithography and surface treatments) have successfully reduced γ_0 by several orders of magnitude (Müller et al., 2019). This relentless engineering of the qubit’s immediate environment is what has enabled the transition into the modern era of multi-qubit processors.

The value of γ_0 is determined by the sum of all channels through which a qubit can interact with its environment. In a solid-state QPU, these mechanisms are numerous and complex. They include capacitive coupling to microscopic two-level system (TLS) defects in the amorphous oxides of the chip, magnetic field noise from nearby nuclear spins in the substrate, and radiative loss through spurious antenna modes in the qubit’s geometry (Müller et al., 2019). Each of these channels contributes to the overall γ_0 . The process of “manufacturing quality” is therefore the process of systematically identifying and eliminating or mitigating these coupling mechanisms through better design, purer materials, and more precise fabrication.

The seven distinct models simulated in this work are architecturally defined by their different initial γ_0 values, representing a spectrum of manufacturing qualities. MODEL_01, the “Idealized Device,” is assigned a γ_0 of $1.0\text{e-}8$, representing a future device fabricated with near-perfect materials. In contrast, MODEL_05, the “Noisy Prototype,” has a γ_0 of $3.0\text{e-}6$, characteristic of a device with significant material defects or poor electromagnetic shielding. The performance difference between these architectural starting points, as will be shown in the analysis, is dramatic, even though all other system parameters are identical.

A common architectural perspective, particularly in classical computing, is that the most important metric for a processor is its scale, i.e., the number of transistors or, in this case, the number of qubits. From this viewpoint, one might argue that it is better to have a large array of lower-quality qubits than a small array of high-quality ones. This “quantity over quality” argument would prioritize scaling up the qubit count over the difficult and expensive process of reducing the initial coupling strength, γ_0 .

This synthesis of QEC theory and the simulation results demonstrates that quality is a prerequisite for quantity to be useful. The threshold theorem of quantum error correction imposes a hard limit on the maximum physical error rate that can be corrected (Google Quantum AI, 2023). A QPU with a poor γ_0 may have a physical error rate so high that QEC is ineffective from the start, regardless of how many qubits it contains. As demonstrated by MODEL_07, if γ_0 is too high, the device is “dead on arrival.” Therefore, the architectural priority must be to first achieve a sufficiently low γ_0 to operate below the fault-tolerance threshold, and only then to scale up the number of qubits.

The intrinsic quality of the QPU, defined by its γ_0 , is only one part of the architectural picture. To function, the QPU must be integrated into a larger system that provides the necessary operating conditions. The most critical of these is the cryogenic environment subsystem, which is responsible for reducing the thermal noise that the QPU is exposed to, thereby minimizing the “T” component of the γT product.

3.2 Cryogenic Environment

Subsystem

The cryogenic environment subsystem, typically a multi-stage dilution refrigerator, is an indispensable architectural component for any scalable quantum computer based on leading solid-state modalities. Its primary function is to reduce the environmental temperature, T , to the millikelvin range, thereby minimizing the population of thermal photons and phonons that can cause decoherence (Krinner et al., 2019). While engineering a lower coupling strength, γ , is a battle fought in the nanofabrication facility, reducing the temperature is a continuous, active battle fought by the cryogenic subsystem during every moment of the computer’s operation. This subsystem is not merely auxiliary equipment; it is an integral and performance-defining part of the quantum computer itself.

The development of quantum computing and the advancement of cryogenic technology have been deeply intertwined. The first experimental qubits of the 1990s were often operated in complex, liquid-helium-based “wet” cryostats. The major breakthrough that enabled the current era of commercial and academic quantum computing was the development of reliable, closed-cycle, “dry” dilution refrigerators. These systems, which can run continuously for months or years without manual intervention, transformed quantum computing from a specialized physics experiment into a stable, 24/7 operational platform (Guan et al., 2025). The ongoing co-evolution of the field sees a demand for ever-increasing cooling power to handle the thermal loads of larger and more complex QPUs.

The mechanism of the cryogenic subsystem is to create a series of progressively colder thermal stages, culminating in a base temperature at the mixing chamber stage that is typically below 20 millikelvin. The QPU is thermally anchored to this stage. The subsystem actively removes heat that leaks into the system from the outside world, primarily through the control and readout wiring, and any heat generated by the operation of the QPU or its associated control electronics. However, the efficiency of this heat removal is limited by the thermal boundary resistance, or Kapitza resistance, at the interface between the QPU chip and its thermal anchor, which can cause the chip to be significantly hotter than the refrigerator’s thermometer indicates (Swartz & Pohl, 1989).

The architectural importance of the cryogenic subsystem is represented in the S3A simulation by the choice of a constant, fixed operating temperature of $T = 0.02$ Kelvin (20 mK). This parameter choice is a direct reflection of the state-of-the-art in dilution refrigerator technology. By holding this temperature constant across all seven simulated models, the analysis implicitly assumes the presence of a perfectly functioning and sufficiently powerful cryogenic subsystem capable of maintaining this stable thermal environment, regardless of the QPU’s properties or the passage of time. This idealization allows the simulation to focus specifically on the effects of the QPU’s intrinsic degradation.

The primary counter-argument to the necessity of this complex cryogenic architecture is the existence of quantum systems that can operate at or near room temperature. Proponents of this view would point to platforms like nitrogen-vacancy centers in diamond as proof that the immense complexity and power consumption of dilution refrigerators are not a fundamental requirement for quantum computation (Doherty et al., 2013). This argument suggests that research efforts should be focused on developing more of these intrinsically stable, room-temperature-capable qubits, which would render the entire cryogenic subsystem obsolete.

While room-temperature platforms are an exciting and important area of research, the synthesis of the current state of the field shows that for scalable, gate-based quantum computing, the leading and most rapidly advancing modalities (superconducting and semiconductor qubits) are fundamentally reliant on cryogenic operation. These platforms offer significant advantages in terms of fast gate speeds and advanced manufacturing capabilities, which are critical for building large-scale processors. The architectural decision to rely on a cryogenic subsystem is therefore a pragmatic trade-off: it accepts the significant engineering complexity of refrigeration in exchange for access to the most promising technologies for achieving large-scale quantum computation in the near term.

The cryogenic subsystem provides the cold, quiet stage upon which the quantum computation can be performed. However, to actually execute an algorithm, a vast and complex layer of classical electronics is required to send instructions to the QPU and interpret the results. This classical control and readout hardware forms the next critical layer of the system architecture, acting as the interface between the classical world of the programmer and the quantum world of the qubits.

3.3 Classical Control

and Readout Hardware

The classical control and readout hardware is the architectural layer that translates abstract algorithmic commands into physical operations on the quantum processor and, conversely, converts the quantum state of the qubits into classical information. This subsystem is a critical performance bottleneck, as the speed, precision, and fidelity of these classical electronics directly limit the overall computational power of the quantum computer (Pauka et al., 2021). A quantum computer is therefore a fundamentally hybrid system, whose performance is co-dependent on the quality of both its quantum and classical components. The architecture must be designed to ensure that the classical hardware is fast and accurate enough to not become the limiting factor.

The evolution of quantum control hardware has been a transition from repurposed laboratory test equipment to highly specialized, integrated systems. Early experiments in the 1990s and 2000s were often controlled by a rack of off-the-shelf arbitrary waveform generators, signal generators, and oscilloscopes, manually synchronized. The need to control larger and larger numbers of qubits with nanosecond timing precision has driven the development of dedicated, multi-channel control platforms (Krinner et al., 2019). The most recent and significant architectural shift is the move to integrate this classical hardware into the cryogenic environment itself, in the form of cryo-CMOS controllers, to overcome the “wiring bottleneck” (Pauka et al., 2021).

The control mechanism involves several stages. A high-level compiler first breaks down a quantum algorithm into a sequence of elementary quantum gates. This sequence is then sent to a pulse-level controller, which translates each gate into a precisely shaped and timed analog signal, typically a microwave pulse. These pulses are generated by arbitrary waveform generators, mixed up to the qubit’s resonant frequency, and sent down transmission lines to the QPU. For readout, a weaker probe signal is sent to the qubit’s resonator, and the reflected or transmitted signal, which is modified by the qubit’s state, is captured. This faint signal is then amplified by a series of low-noise amplifiers (often including a quantum-limited amplifier at the coldest stage) and finally digitized and processed to determine the classical ‘0’ or ‘1’ result (Pauka et al., 2021).

The capabilities of this classical control architecture are represented in the S3A simulation by the parameter t_{gate} , the gate operation time, which is set to a fixed value of 4 nanoseconds. This choice reflects the performance of a state-of-the-art classical control system capable of generating the very fast pulses required to operate a 5 GHz superconducting qubit. This fixed parameter implicitly assumes that the classical hardware is perfectly stable and does not degrade over time, an idealization that allows the simulation to focus solely on the degradation of the quantum components. The speed of this hardware is critical; a faster t_{gate} would lead to a lower physical error rate, ϵ_p , for the same coherence time.

A counter-argument from a purely quantum-focused perspective might be that the classical control hardware is a secondary concern. This view would posit that as long as the quantum coherence is long enough, the speed of the classical gates is not a primary issue. If a qubit can hold its state for seconds, it doesn’t matter if a gate takes microseconds to perform. This line of reasoning would prioritize investment in improving qubit coherence above all else, assuming that the classical control problem is a straightforward engineering task that can be solved later.

The synthesis of different hardware platforms reveals that a balance between quantum coherence and classical control speed is essential. The case of trapped ions provides a perfect illustration: they possess extraordinarily long coherence times (minutes) but are limited by relatively slow gate speeds (microseconds) (Bruzewicz et al., 2019). This makes it challenging for them to execute very deep circuits with many sequential operations, a problem known as the “depth-to-coherence ratio.” A viable system architecture requires both a long coherence time *and* fast gates. The classical control hardware is therefore not a secondary concern but a co-equal partner in the architectural challenge of maximizing the number of useful operations that can be performed before decoherence.

The classical control and readout hardware forms the interface for executing operations and measuring their outcomes. However, in a fault-tolerant architecture, these raw measurement outcomes are not the final result. They are, instead, error syndromes that must be fed into another, even more specialized classical computational system: the real-time quantum error correction decoder engine. This component is the brain of the operation, responsible for interpreting the noisy data from the QPU and orchestrating its correction.

3.4 Real-Time QEC Decoder

Engine

The real-time quantum error correction (QEC) decoder engine is a specialized, high-performance classical computer that forms the logical core of a fault-tolerant quantum architecture. Its sole function is to process the continuous stream of noisy error syndrome measurements from the QPU, deduce the most likely physical errors that have occurred, and dispatch corrective operations, all within a fraction of the qubit coherence time (Sivak et al., 2023). This engine is the component that actively implements the error correction, and its computational capacity represents a hard architectural limit on the level of noise and degradation that the system can tolerate.

The historical development of QEC decoders has been a transition from a purely theoretical concept to a critical piece of experimental infrastructure. For many years, decoding was an offline, post-processing step performed on a conventional computer long after the quantum experiment was finished (Google Quantum AI, 2023). The breakthrough that enabled modern QEC experiments was the development of real-time decoders, typically implemented on FPGAs (Field-Programmable Gate Arrays), that could perform the decoding calculations fast enough to keep pace with the quantum hardware. The ongoing challenge is to design decoders that are not only fast but also scalable to the millions of qubits that a future fault-tolerant computer will require.

The architectural mechanism of the decoder engine is a tight feedback loop. In each cycle of the QEC protocol, the readout hardware delivers a new set of syndrome bits to the decoder. These bits do not indicate where the errors are, but rather where the “boundaries” of the error chains are. The decoder’s task is to solve a complex inverse problem: given this boundary information, what is the smallest and most probable set of physical errors that could have produced it? For the surface code, this problem can be mapped onto a well-known problem in graph theory called Minimum Weight Perfect Matching. The decoder runs this algorithm on the syndrome graph, and the solution dictates which physical qubits need to have corrective Pauli operations (X, Y, or Z gates) applied to them (Sivak et al., 2023).

The computational load placed on this decoder engine is the quantity represented by the “Resource Cost,” R_{cost} , in the S3A simulation. The model’s use of $R_{\text{cost}} = d_{\text{practical}}^2$ is a proxy for the escalating complexity of the decoding problem as the code distance increases. As the physical hardware degrades and $\epsilon_p(t)$ rises, the system must increase $d_{\text{practical}}$ to compensate. The simulation logs show this directly: for the “Noisy Prototype” (MODEL_05), R_{cost} triples over 14 months, from 289 to 961 arbitrary units. This represents the exponentially increasing demand placed on the decoder engine as it struggles to correct the errors of the aging QPU.

A counter-argument might be that the decoding can be done in a non-real-time or “offline” manner. In this architectural model, one would simply store the history of all syndrome measurements during the quantum algorithm’s execution. After the algorithm is finished, a powerful classical supercomputer could take its time to perform the decoding and reconstruct the final, corrected logical state. This would remove the stringent low-latency requirement and potentially allow for more sophisticated but slower decoding algorithms.

While offline decoding is a useful tool for analyzing experimental data, it is insufficient for scalable, fault-tolerant computation. The core principle of fault tolerance is to correct errors as they happen, *before* they have a chance to propagate and spread through subsequent two-qubit gates, where they can become complex, correlated logical

errors. If error correction is not performed in real-time, the accumulation of uncorrected errors would quickly overwhelm the code's ability to correct them, rendering the computation useless. The synthesis of QEC theory and experimental practice, as demonstrated in the latest real-time experiments, confirms that a low-latency feedback loop is a non-negotiable architectural requirement (Sivak et al., 2023).

The complete system architecture can now be understood as a tightly integrated, hybrid quantum-classical machine operating in a continuous, dynamic feedback cycle. The state of the quantum hardware dictates the demands on the classical decoder, and the capacity of the decoder dictates the ultimate viability of the entire system. This dynamic interplay, which unfolds over the operational lifetime of the device, is the degradation-feedback loop.

3.5 Degradation-Feedback Loop

The complete quantum computing system architecture operates within a continuous, long-term degradation-feedback loop. This loop is the central dynamic that governs the device's operational lifecycle. It begins with the slow, inevitable physical degradation of the quantum processing unit, which in turn forces an escalating, resource-intensive response from the classical quantum error correction subsystem. The viability of the entire architecture is determined by its ability to sustain this feedback loop until either the physical hardware becomes uncorrectable or the classical resources are exhausted.

This architectural concept of a degradation-feedback loop provides a long-term temporal context for quantum computing that is often missing from performance benchmarks, which typically represent a static snapshot of a device at its peak. By framing the system's operation as a continuous, evolving process, this perspective aligns the analysis of quantum computers with the established fields of reliability engineering and lifecycle management used for other high-performance, critical systems like satellites or supercomputers. It acknowledges that the challenge is not just to build a machine that works, but to build a machine that can continue to work for a useful period of time.

The feedback loop proceeds through a clear, causal five-step cycle that repeats over the device's lifetime. First, the physical QPU degrades, causing its mean system-environment coupling, $\gamma(t)$, to slowly increase. Second, this increased coupling leads to a higher measured physical error rate, $\epsilon_p(t)$. Third, the classical control system detects this increase and, to maintain the target logical error rate, determines that a stronger QEC protocol is needed, thus increasing the required code distance, $d(t)$. Fourth, this higher code distance increases the computational load, $R_{\text{cost}}(t)$, on the real-time QEC decoder engine and the number of physical qubits that must be actively managed. Fifth, the loop repeats, with the system continuing to operate at this higher resource cost until further degradation forces the next escalation.

The temporal progression observed in the S3A simulation logs for each of the seven models is a direct, numerical simulation of this degradation-feedback loop. For example, the log for the "Workhorse Device" (MODEL_03) provides a clear narrative of this cycle in action. At $t=187.5$ days, the accumulated degradation of $\gamma(t)$ causes $\epsilon_p(t)$ to cross a new threshold, forcing the system to increase its code distance from 11 to 13. This, in turn, causes the resource cost to jump from 121 to 169 units. The system then operates in this new, more expensive state until $t=395.8$ days, when the loop triggers again, pushing the code distance to 15 and the resource cost to 225.

A counter-argument could be that this feedback loop is not necessary and that a simpler, static architecture would suffice. In this view, a system should be designed from the outset with a fixed, maximum code distance sufficient to handle the predicted end-of-life error rate. The system would operate in this high-resource-cost mode from day one. This would eliminate the need for adaptive logic and would result in a system whose performance is static and predictable, albeit inefficient in its early life.

While a static, worst-case design is architecturally simpler, it is profoundly inefficient. The synthesis of the simulation results shows that the resource cost escalates dramatically over time. Designing a system to handle the end-of-life resource cost from day one would mean massively overprovisioning the classical control hardware, which would sit mostly idle for a significant portion of the device's life. A dynamic, adaptive architecture, as described by the feedback loop, allows for a much more efficient allocation of resources, where the computational cost is always matched to the current, real-world needs of the physical hardware. This “just-in-time” approach to error correction is far more practical and economical.

The existence of this degradation-feedback loop forces a critical, high-level strategic decision in the design of any quantum computing architecture. It creates a fundamental trade-off between investing in the initial quality of the quantum hardware versus investing in the power and capacity of the classical hardware that must compensate for its flaws. This trade-off between physical purity and computational redundancy is a defining feature of the architectural landscape.

3.6 Engineering

Trade-offs: Purity vs. Redundancy

A core architectural trade-off in the design of a fault-tolerant quantum computer is the strategic allocation of resources between two competing philosophies: “purity” and “redundancy.” The purity strategy prioritizes investment in fundamental materials science and fabrication techniques to create a quantum processor with the lowest possible initial coupling strength, γ_0 (Burkard et al., 2021). The redundancy strategy, in contrast, accepts a higher initial coupling and instead invests in a more powerful and expensive classical control system capable of managing the higher quantum error correction overhead that will result (Google Quantum AI, 2023). This trade-off represents a fundamental economic and engineering decision that shapes the entire system architecture.

This purity-versus-redundancy trade-off is not unique to quantum computing, but its implications are particularly extreme in this domain. In the history of classical computing, a similar tension existed between building more reliable vacuum tubes versus developing error-correcting codes to tolerate their failures. The eventual dominance of the highly reliable transistor largely favored the purity approach. In quantum computing, however, where perfect physical reliability is impossible, the optimal balance is far less clear. The field is currently exploring both paths simultaneously: some research groups focus on achieving new records in coherence times (purity), while others focus on scaling up qubit numbers and implementing more powerful QEC codes (redundancy).

The mechanism of the trade-off is directly visible in the simulation model. A lower initial γ_0 , representing the purity approach, results in a lower physical error rate, ϵ_p , for a longer period. This means the system can operate with a lower QEC code distance, d , leading to a significantly lower sustained resource cost, R_{cost} . This path requires a higher upfront investment in research and development and advanced fabrication facilities. Conversely, a higher γ_0 , representing the redundancy approach, may be cheaper and faster to manufacture, but it immediately requires a higher code distance and a more powerful classical decoder engine to function. This path shifts the cost from the fabrication facility to the classical control infrastructure and accepts a shorter operational lifetime.

The S3A simulation logs provide a stark, quantitative illustration of this trade-off. Comparing the “Idealized Device” (MODEL_01) with the “Noisy Prototype” (MODEL_05) is instructive. The idealized device, with its excellent γ_0 , operates for over 14 months with a resource cost that only rises from 49 to 81 arbitrary units. The noisy prototype, to perform the exact same logical computation, requires a resource cost that starts at 289 units and escalates to 961 units over the same period. This demonstrates that an initial factor of 300 improvement in the purity of the QPU (i.e., a lower γ_0) results in a more than tenfold reduction in the sustained operational cost.

A counter-argument, particularly from a near-term commercial perspective, might be that one should always choose the path of redundancy. This view would hold that it is faster and more predictable to scale up classical computing power—a well-understood technology—than it is to wait for uncertain breakthroughs in fundamental materials science. This “build it now with what we have” approach would favor using currently available, moderately noisy qubits and compensating with a massive investment in classical control hardware, with the goal of reaching a useful scale as quickly as possible.

The synthesis of the simulation results suggests that while the redundancy approach is viable up to a point, the purity of the underlying hardware is the ultimate long-term driver of efficiency and scalability. The quadratic scaling of the resource cost with the code distance means that there is a point of diminishing returns for the redundancy strategy. As shown by the “Critical Threshold Device” (MODEL_06), a device with a sufficiently poor γ_0 faces such a rapid explosion in resource costs that no practical classical computer could keep pace. Therefore, the optimal architecture likely involves a balanced approach: investing in purity to ensure the physical error rate is well below the fault-tolerance threshold, and then applying the necessary level of redundancy to achieve the target logical error rate.

This fundamental trade-off between purity and redundancy directly informs the ultimate architectural goal: to design a system that can operate for a useful lifetime without succumbing to failure. The entire architectural framework—from the QPU to the cryogenic subsystem to the classical decoder—must be designed with a clear and quantitative understanding of the conditions that define the end of this operational lifetime.

3.7 Defining

System Lifetime and Failure Conditions

The operational lifetime of the quantum computing system architecture is formally defined as the duration of time, starting from its initial deployment, until it breaches one of two critical and non-recoverable failure conditions. This provides a concrete, quantitative metric for the useful life of the device, moving beyond simple performance benchmarks to a holistic assessment of its long-term reliability and sustainability. The architecture is considered to have failed when it can no longer guarantee the execution of a fault-tolerant algorithm at the target logical error rate.

This lifecycle-based definition of success and failure is standard practice in mature engineering disciplines. For a commercial aircraft, the lifetime is defined by a certain number of flight hours or pressurization cycles, after which material fatigue makes it unsafe to operate. For a satellite, the lifetime is often determined by the amount of propellant available for station-keeping. By applying a similar, formally defined end-of-life concept to a quantum computer, the architectural analysis shifts from a purely scientific endeavor to a rigorous engineering one, focused on reliability, longevity, and the total cost of ownership.

The two failure conditions that define the system’s end-of-life are architecturally distinct. The first is a **physical failure**. This occurs when the physical error rate, $\epsilon_p(t)$, of the degrading QPU rises above the fault-tolerance threshold, ϵ_{th} , of the chosen quantum error correction code. At this point, the QEC algorithm becomes counterproductive, introducing more errors than it corrects, and fault-tolerant computation is no longer possible (Google Quantum AI, 2023). The second is a **resource failure**. This occurs when the required resource cost of QEC, $R_{cost}(t)$, exceeds the designed maximum capacity of the classical real-time decoder engine. At this point, the classical hardware can no longer keep up with the demands of the quantum processor, leading to a cascade of uncorrected errors.

The S3A simulation suite was explicitly designed to capture these two failure modes. The simulation of the “Sub-Threshold Device” (MODEL_07) provides a clear example of a physical failure. Its initial coupling, γ_0 , was so high that its ϵ_p at $t=0$ was already above the ϵ_{th} of 0.01, causing the simulation to terminate and report an immediate, non-recoverable failure. In contrast, the other viable models, such as the “Critical Threshold Device”

(MODEL_06), illustrate the trajectory towards a resource failure. While its ϵ_p remains below the threshold, its R_{cost} grows exponentially, indicating that it would inevitably surpass the capacity of any finitely provisioned classical control system.

A potential counter-argument is that a device's life could be extended indefinitely through periodic recalibration or minor repairs. This view would suggest that "end-of-life" is not a fixed point, but that the system can be maintained and serviced to keep it within operational parameters, much like a classical car can be kept running with regular maintenance. From this perspective, the degradation process is not a one-way street, and the failure conditions are not necessarily permanent.

While frequent recalibration is an essential part of operating any quantum computer and can compensate for short-term drifts, it cannot reverse the fundamental, long-term material degradation that drives the increase in the baseline coupling strength, γ . The accumulation of microscopic defects in the solid-state substrate is, for all practical purposes, irreversible. The synthesis of the model and the physical understanding of material aging confirms that the architecture must be designed with a finite operational lifetime in mind. The goal of the architecture is to make this lifetime as long and as productive as possible, but it cannot be infinite.

With the complete system architecture and its operational lifecycle now fully defined, the final step is to analyze the quantitative results of the simulation. The following analysis section will systematically walk through the simulated lifecycles of each of the seven device classes, from the idealized to the impossible, drawing concrete conclusions about the relationship between initial device quality and long-term computational viability.

4.0 Analysis

4.1 Idealized Device Lifecycle

An idealized quantum device, characterized by an exceptionally low initial system-environment coupling strength, is projected to maintain computational viability with only a minimal escalation in resource cost over a simulated 14-month operational period. This result demonstrates that achieving superior initial material quality and qubit isolation is the most effective strategy for ensuring a long and computationally efficient device lifetime. The stability of this idealized model serves as a crucial benchmark, representing the ultimate goal of hardware development and quantifying the profound long-term benefits of investing in fundamental materials science and fabrication precision.

This model represents a future, highly advanced quantum processor that surpasses the quality of current state-of-the-art laboratory results. It is defined by an initial coupling strength of $\gamma_0 = 1.0\text{e-}8$. While not yet achieved in practice for scalable platforms like transmons, this level of isolation is theoretically possible and serves as a target for next-generation device engineering. The analysis of this model's lifecycle is therefore not a characterization of a current device, but a forward-looking projection of the potential return on investment for research into novel materials and defect-reduction techniques.

The mechanism behind the device's exceptional stability is the large initial gap between its physical error rate and the fault-tolerance threshold. Due to the extremely low γ_0 , the initial physical error rate, ϵ_p , is very small. The slow, exponential degradation of $\gamma(t)$ therefore proceeds for a long time before $\epsilon_p(t)$ crosses the boundary that necessitates an increase in the quantum error correction code distance. This results in a long period of stable, low-cost operation, followed by infrequent, small steps in resource allocation. The system spends the vast majority of its life in a highly efficient operational state.

The numerical logs for MODEL_01 provide clear, quantitative evidence of this behavior. At $t=0$, the system requires a minimal quantum error correction code distance of $d=7$, corresponding to a resource cost of $R_{\text{cost}} = 49$ arbitrary units. This cost remains perfectly stable for the first three months of operation. It is not until $t=104.2$ days that the

accumulated degradation forces a single increase in the code distance to $d=9$. This raises the resource cost to 81 units, where it then remains for the rest of the 14-month simulation. The total increase in operational cost over the entire period is a mere 65%.

A skeptical counter-argument might be that such a low level of degradation is unrealistic and that all devices, regardless of their initial quality, will degrade at a much faster rate when subjected to the stresses of continuous operation. This perspective would suggest that the idealized model is a fantasy and that its projected stability is an artifact of an overly optimistic degradation constant (k_{degrade}). From this viewpoint, the long, stable lifetime is a simulation artifact, not a realistic engineering target.

While the exact degradation rate of future devices is an open question, the synthesis of the model's logic demonstrates a fundamental principle: a lower starting point in a race against exponential growth yields disproportionately large benefits. By starting with a γ_0 that is 20 times better than the “Workhorse” device, the idealized model gains a significant temporal buffer before the escalating error rate becomes problematic. The simulation's result—a long and efficient operational lifetime—is a direct and robust consequence of this initial advantage. This validates the architectural strategy of prioritizing fundamental materials science to achieve the lowest possible initial coupling strength.

The exceptional performance of this idealized device provides a powerful aspirational benchmark. However, to ground the analysis in the present, it is necessary to compare this future projection against the lifecycle of a device that represents the current state of the art. The “Hero Device” model serves this purpose, providing a more realistic, albeit still optimistic, assessment of a top-tier contemporary quantum processor.

4.2 Hero Device Lifecycle

A “hero” device, representing the performance of a state-of-the-art, record-setting laboratory prototype, is projected to exhibit a stable operational life for a significant period but shows the clear onset of resource cost escalation within its first year. This analysis indicates that even the best currently achievable devices are subject to measurable degradation on a medium-term timescale. The lifecycle of this model demonstrates that system architectures must be designed with the expectation of adapting to a rising error rate, even when using the highest quality qubits available today.

This model is defined by an initial coupling strength of $\gamma_0 = 5.0e-8$. This value is chosen to be representative of a top-tier experimental result that might be reported in a leading scientific journal—a device that has been meticulously fabricated and selected for its superior performance, but which may not be representative of the average device in a larger array. The analysis of its lifecycle provides insight into the upper bound of performance that can be expected from the current generation of quantum hardware technology.

The underlying mechanism of the hero device's lifecycle is similar to that of the idealized device, but accelerated. With a γ_0 that is five times higher, the starting physical error rate ϵ_p is closer to the next quantum error correction threshold. As $\gamma(t)$ degrades exponentially, it crosses this threshold much sooner than in the idealized case. This forces the system to allocate more classical resources to error correction at an earlier stage in its operational life, marking the beginning of the escalating cost curve that characterizes the device's long-term trajectory.

The numerical logs for MODEL_02 clearly illustrate this accelerated timeline. The device begins at $t=0$ with a required code distance of $d=9$, corresponding to a resource cost of $R_{\text{cost}} = 81$ units. It maintains this stable and efficient operational state for a considerable period. However, at $t=208.3$ days, approximately seven months into its life, the accumulated degradation forces an increase in the code distance to $d=11$. This results in a significant 49% jump in the resource cost to 121 units, a level that is then maintained for the remainder of the simulation.

A counter-argument could be that a single “hero” device is not a meaningful subject for a lifecycle analysis. Such devices are often operated for short periods to achieve a specific benchmark and are not intended for long-term, stable operation. From this perspective, analyzing the degradation of a device that is not representative of a commercial-grade, production system is an academic exercise with little practical relevance for the future of scalable quantum computing.

This analysis, however, serves a crucial purpose. By demonstrating that even the highest-quality, most carefully engineered devices of the current era are subject to significant degradation within months, it provides a powerful argument against the notion that the problem of device stability has been solved. The synthesis of this result with the broader understanding of material science confirms that there is no “silver bullet” qubit that is immune to aging. The lifecycle simulation of the hero device establishes a realistic, data-driven upper bound on the expected stability of current technology and reinforces the necessity of designing adaptive, fault-tolerant systems.

While the hero device represents the pinnacle of current laboratory achievement, it does not reflect the performance of the average, statistically typical qubit that will form the backbone of a large-scale quantum computer. To understand the operational reality for such a machine, it is necessary to analyze the lifecycle of the “Workhorse Device,” which is defined by a more conservative and representative initial coupling strength.

4.3 Workhorse Device Lifecycle

A “workhorse” device, defined by an average and statistically representative initial coupling strength, demonstrates a clear and continuous lifecycle of escalating resource costs. The analysis of this model shows that a typical production-grade quantum processor must be supported by a classical control system capable of dynamically increasing its computational resource allocation to quantum error correction multiple times throughout its operational life. This model most accurately reflects the central thesis of this work: that maintaining computational viability is an active and increasingly costly process of compensating for hardware degradation.

This model is defined by an initial coupling strength of $\gamma_0 = 2.0\text{e-}7$. This value is chosen to represent the mean or median performance of qubits in a large, wafer-scale fabrication run. While a few “hero” devices on the wafer might have a lower γ_0 , and many others will be worse, this value represents the typical quality that a scalable architecture must be designed to work with. The analysis of this model’s lifecycle is therefore the most relevant for predicting the behavior and operational costs of near-term, large-scale quantum computers.

The mechanism driving the workhorse lifecycle is the relentless progression of the degradation-feedback loop. Starting with a moderate γ_0 , the device’s physical error rate $\epsilon_p(t)$ begins at a level that already requires a significant quantum error correction overhead. As $\gamma(t)$ degrades, $\epsilon_p(t)$ steadily worsens, forcing the system to repeatedly climb the ladder of quantum error correction code distances. Each step up this ladder corresponds to a discrete, significant jump in the computational resources required to keep the logical error rate stable.

The numerical logs for MODEL_03 provide a clear and compelling narrative of this escalating cost. The device begins its life at $t=0$ with a required code distance of $d=11$, corresponding to a resource cost of $R_{\text{cost}} = 121$ units. After approximately six months ($t=187.5$ days), the first resource escalation occurs, pushing the code distance to $d=13$ and the cost to 169 units. The system operates in this more expensive state for another seven months until, at $t=395.8$ days, a second escalation is required, increasing the distance to $d=15$ and the cost to 225 units. Over the 14-month simulation, the operational cost of the device nearly doubles, increasing by 86%.

A counter-argument might be that the observed increase in resource cost is not a fundamental problem but a manageable operational expense. From a systems-engineering perspective, as long as the classical control system is designed with sufficient headroom to accommodate this increase, the device remains perfectly viable. This view

would frame the escalating cost not as a looming failure, but simply as a predictable and budgetable aspect of the system's total cost of ownership.

While it is true that the system remains technically viable throughout the simulation, the analysis highlights a critical architectural and economic challenge. The fact that the operational cost of the quantum computer is not fixed, but is a steeply increasing function of time, has profound implications. It means that the classical control system must be significantly overprovisioned from day one, with expensive computational resources sitting idle for the first year of the device's life. This synthesis reveals that the "cost" of degradation is paid not just at the end of life, but throughout the entire operational period in the form of underutilized capital investment.

The lifecycle of the workhorse device, with its doubling of operational cost, represents the expected behavior of a typical, modern quantum processor. To provide a more complete picture, it is instructive to compare this to the lifecycle of an older, previous-generation device, which would be characterized by a significantly higher initial coupling strength and an even more aggressive cost escalation curve.

4.4 Legacy Device Lifecycle

A legacy device, representing an older generation of quantum hardware with a higher initial coupling strength, is projected to exhibit a rapid and continuous escalation of quantum error correction resource costs from the very beginning of its lifecycle. The analysis of this model demonstrates that lower-quality initial hardware is not only less efficient but also has a much steeper trajectory toward resource exhaustion. This validates the historical progression of the field, showing that each generational improvement in reducing γ_0 yields significant, compounding benefits in long-term operational sustainability.

This model is defined by an initial coupling strength of $\gamma_0 = 8.0e-7$, a value four times higher than that of the "Workhorse" device. This is chosen to be representative of the quality of devices that were considered state-of-the-art several years ago. By simulating the lifecycle of this "Legacy Device" under the same conditions as the modern hardware, the analysis can quantify the practical impact of the technological progress that has been made in materials science and fabrication over the past hardware generation.

The mechanism driving the legacy device's rapid cost escalation is its high initial physical error rate. Because γ_0 is already large, the device starts its life with a ϵ_p that is much closer to the subsequent quantum error correction thresholds. This means that even a small amount of absolute degradation in $\gamma(t)$ is sufficient to trigger the need for a higher code distance. The device is therefore forced to climb the ladder of QEC resource costs much more quickly and frequently than a higher-quality, modern device.

The numerical logs for MODEL_04 clearly illustrate this steep cost curve. The device begins operation at $t=0$ already requiring a high code distance of $d=13$, corresponding to a resource cost of $R_{\text{cost}} = 169$ units. The first resource escalation occurs after just 83.3 days, pushing the code distance to $d=15$ and the cost to 225 units. The system requires two further escalations during the 14-month period, to $d=17$ at 229.2 days and $d=19$ at 354.2 days. By the end of the simulation, the resource cost has reached 361 units, representing a 114% increase from its already high starting point.

A counter-argument could be that such legacy devices can still be useful for less demanding tasks that do not require full fault tolerance. For example, they could be used for experiments in the NISQ paradigm or for educational purposes, where a higher logical error rate is acceptable. From this perspective, the escalating cost of *fault-tolerant* operation is irrelevant if the device is repurposed for a different class of computation where such stringent error correction is not required.

This is a valid point, and it highlights the importance of matching the hardware to the computational task. However, within the context of building a scalable, fault-tolerant quantum computer capable of solving classically intractable problems, the analysis holds. The simulation demonstrates that the legacy device is economically and computationally inefficient for this purpose. The synthesis of the results shows that the total computational effort (the integral of R_{cost} over time) expended by the legacy device is far greater than that of the workhorse device for the same logical task. This confirms that using older, noisier hardware for fault-tolerant computation is a fundamentally inefficient strategy.

The steep cost curve of the legacy device illustrates the challenges of working with moderately noisy hardware. To understand the absolute limits of viability, it is necessary to analyze the lifecycle of a device that is even noisier from its inception. The “Noisy Prototype” model is designed to probe this boundary, where the initial resource requirements are extreme and their rate of escalation is even more severe.

4.5 Noisy Prototype Lifecycle

A noisy prototype device, characterized by a high initial coupling strength, is shown to be technically viable but practically unsustainable for fault-tolerant operation due to its extremely high and rapidly escalating resource costs. The analysis of this model demonstrates that there is a point where the overhead of quantum error correction becomes so large that it renders the system impractical, even if it has not yet breached a hard physical failure threshold. This highlights the critical role of economic and computational efficiency in the definition of a useful quantum computer.

This model is defined by a high initial coupling strength of $\gamma_0 = 3.0\text{e-}6$. This value is representative of an early-stage experimental device, perhaps one that is poorly shielded from magnetic noise, fabricated with a new and unoptimized process, or designed to test a novel but noisy qubit architecture. The purpose of analyzing its lifecycle is to understand the behavior of a system that is operating close to the edge of what is correctable, and to quantify the extreme measures required to maintain its logical integrity.

The mechanism at play in this model is the same degradation-feedback loop as in the previous cases, but operating in a much more extreme regime. The high initial γ_0 forces the system to begin its life with a very high required code distance just to meet the target logical error rate. Because the physical error rate is already in a sensitive part of the error-suppression curve, even small absolute increases in $\gamma(t)$ due to degradation cause a large relative increase in $\epsilon_p(t)$. This forces the system to make frequent and large jumps in the required code distance, leading to an almost exponential growth in the resource cost.

The numerical logs for MODEL_05 provide a stark picture of this unsustainable cost explosion. The device begins operation at $t=0$ already requiring a code distance of $d=17$, corresponding to a resource cost of $R_{\text{cost}} = 289$ units. Over the course of the 14-month simulation, the system is forced to increase its code distance eight separate times, eventually reaching $d=31$. This causes the resource cost to more than triple, skyrocketing to 961 units by the end of the period. The system spends more computational effort on correcting its own errors than on performing any useful logical computation.

A purely theoretical counter-argument might be that as long as the physical error rate remains below the fault-tolerance threshold, and as long as one has access to a sufficiently powerful classical computer (perhaps a hypothetical, infinitely powerful one), the device is still viable. From this perspective, a high resource cost is merely an engineering detail, not a fundamental barrier to viability. If the goal is simply to prove that fault tolerance is possible, then the cost of achieving it is irrelevant.

This synthesis, however, must be grounded in the reality of building a functional and useful machine. A quantum computer is not a theoretical construct; it is a physical artifact that must be built and operated within finite economic and computational budgets. The analysis of the noisy prototype demonstrates that there is a clear distinction between what is theoretically possible and what is practically feasible. A system that requires its operational cost to triple in just over a year is not a sustainable architecture. It proves that simply “being correctable” is not a sufficient condition for a device to be considered a viable candidate for scalable quantum computing.

The lifecycle of the noisy prototype brings the analysis to the very edge of practical viability. It begs the question: what happens when a device is manufactured with an initial quality that is just slightly worse? The “Critical Threshold Device” model is designed to explore this precise boundary, revealing the behavior of a system teetering on the precipice of failure from its very first day of operation.

4.6 Critical Threshold

Device Lifecycle

A device manufactured at the critical threshold of viability, with the highest possible initial coupling strength that is still theoretically correctable, exhibits an extreme and immediate explosion in resource cost. The analysis of this model reveals a phase transition in the system’s behavior, where the operational cost becomes so high and escalates so rapidly that the device is rendered practically unsustainable. This demonstrates that the boundary of viability is not a gentle slope but a sharp cliff, and that devices operating too close to this edge face a fundamentally different and more challenging lifecycle.

This model is defined by an initial coupling strength of $\gamma_0 = 4.9\text{e-}6$. This specific value was carefully chosen through preliminary simulations to be just below the point where the initial physical error rate would exceed the fault-tolerance threshold of $\epsilon_{th} = 0.01$. It therefore represents the absolute worst-case scenario for a device that is still functional at $t=0$. The analysis of its lifecycle is a stress test of the quantum error correction system in its most challenging operational regime.

The mechanism driving this model’s behavior is the extreme non-linearity of the quantum error correction scaling formula when the physical error rate ϵ_p is very close to the threshold ϵ_{th} . In the formula for the required code distance, $d \propto 1 / \ln(\epsilon_p / \epsilon_{th})$, the denominator approaches zero as ϵ_p approaches ϵ_{th} . This causes the required d to diverge, increasing dramatically with even the slightest worsening of the physical error rate. The device is therefore trapped in a state of extreme sensitivity, where the smallest amount of material degradation forces a massive and disproportionate response from the QEC system.

The numerical logs for MODEL_06 provide a dramatic illustration of this resource cost explosion. The device begins at $t=0$ with a very high required code distance of $d=19$, corresponding to a resource cost of $R_{\text{cost}} = 361$ units. Unlike the other models which experience periods of stability, this device is forced to increase its code distance at nearly every 20-day time step. Over the 14-month simulation, the code distance is forced to increase eleven times, reaching a staggering $d=41$. This causes the resource cost to skyrocket to 1681 units, an increase of 460% from its already high starting point.

A counter-argument could be that if the system is not technically failing—that is, if ϵ_p remains below ϵ_{th} —then it is still working. From a purely definitional standpoint, as long as the resource cost is finite, the system is viable. This view would hold that the extreme cost is an engineering problem to be solved with a sufficiently powerful classical computer, but it does not represent a fundamental failure of the quantum architecture itself.

This synthesis of the simulation results and practical engineering constraints leads to a more nuanced definition of failure. While the device has not breached the *physical* failure condition, it has clearly entered a state of *resource-based failure*. The exponential-like growth in its operational cost demonstrates a clear phase transition in its

behavior. No realistically designed classical control system could be provisioned to handle a nearly 5x increase in its computational load over a single year. The analysis of this model proves that there is a practical boundary of viability that is reached long before the hard physical limit of the threshold theorem.

The critical threshold device demonstrates the behavior of a system at the absolute limit of what is correctable. This naturally leads to the final question: what happens if a device is manufactured with an initial quality that falls on the other side of this sharp boundary? The final model, the “Sub-Threshold Device,” is designed to provide a definitive answer and conclude the analysis by demonstrating a hard, indisputable failure condition.

4.7 Sub-Threshold Device

Failure

A device manufactured with an initial system-environment coupling strength that places its physical error rate above the fault-tolerance threshold is non-viable from the moment of its creation. The analysis of this model demonstrates the existence of a hard, physical limit to the efficacy of quantum error correction. It proves that QEC is not a magical solution that can fix any error rate; it is a powerful tool that works only when the underlying physical hardware is “good enough.” This finding establishes a clear, non-negotiable target for hardware developers: achieving a physical error rate below the threshold is the absolute, primary requirement for building a fault-tolerant quantum computer (Google Quantum AI, 2023).

This model is defined by an initial coupling strength of $\gamma_0 = 6.0\text{e-}6$. This value is chosen to be just slightly higher than that of the “Critical Threshold Device,” placing it definitively on the wrong side of the viability boundary. This represents a device that has failed its initial quality control check—a result of a poor fabrication run, a critical design flaw, or severe material contamination. The analysis of this model is not a lifecycle simulation, but a static, $t=0$ proof of a fundamental failure mode.

The mechanism of failure is a direct consequence of the threshold theorem of quantum error correction. The theorem, which is the foundation of all fault-tolerant theory, states that for a QEC code to be effective, the physical error rate ϵ_p must be below a certain threshold value, ϵ_{th} (Google Quantum AI, 2023). If ϵ_p is above this threshold, the process of measuring the error syndromes and applying corrections will, on average, introduce more new errors into the system than it fixes. The QEC protocol becomes actively harmful, amplifying noise rather than suppressing it, and the logical error rate will be even higher than the physical error rate.

The numerical log for MODEL_07 provides the most succinct and definitive result of the entire simulation suite. At the very first time step, $t=0$, the simulation calculates the initial physical error rate based on $\gamma_0 = 6.0\text{e-}6$. This ϵ_p is found to be greater than the hard-coded fault-tolerance threshold of $\epsilon_{th} = 0.01$. The simulation immediately terminates and prints the explicit log message: “STATUS: FAILURE. Physical error rate exceeds threshold.” This is not a projection of a future failure; it is a declaration of a present and irreversible condition.

Given the definitive nature of the threshold theorem, there is no credible scientific counter-argument to this result. One could hypothetically argue that a different QEC code with a higher threshold might be able to correct this device. For example, if a new code were invented with a threshold of $\epsilon_{th} = 0.02$, this device might become viable. This, however, is not a counter-argument to the failure of the current system, but a proposal for a different, hypothetical system.

The synthesis of this result with the established theory of fault tolerance is absolute. The failure of the sub-threshold device is a hard, physical limit. It demonstrates that there is a clear, bright-line distinction between a merely “noisy” device and a “non-functional” one (Google Quantum AI, 2023). This finding has profound architectural

implications: the primary goal of hardware engineering must be to manufacture qubits with a γ_0 that places them comfortably below the fault-tolerance threshold. No amount of classical processing power or cleverness in the QEC decoder can compensate for a failure to meet this fundamental physical requirement.

The definitive failure of the sub-threshold device concludes the analysis of the computational matrix. The simulation suite has successfully mapped the entire spectrum of device viability, from the long-term stability of an idealized, future device to the immediate, uncorrectable failure of a low-quality one. This comprehensive analysis, which has quantitatively demonstrated the race between material degradation and the escalating cost of error correction, provides a robust and physically grounded framework for understanding the long-term challenges and strategic imperatives in the development of scalable, fault-tolerant quantum computers.

References

- Aaronson, S. (2013). **Quantum Computing and the Ultimate Limits of Computation**. Computing Research Association.
- Abobeih, M. H., Wang, Y., Randall, J., Loenen, S. J., Bradley, C. E., Taminiau, T. H., Degen, C. L., Markham, M., Twitchen, D. J., & Cramer, J. (2022). Fault-tolerant operation of a logical qubit in a diamond quantum processor. *Nature*, 606(7914), 507–511. <https://doi.org/10.1038/s41586-022-04819-9>
- Brandão, F. G. S. L., Horodecki, M., Oppenheim, J., Renes, J. M., & Spekkens, R. W. (2013). Resource Theory of Quantum States Out of Thermal Equilibrium. *Physical Review Letters*, 111(25), 1. <https://doi.org/10.1103/PhysRevLett.111.250404>
- Breuer, H. P., & Petruccione, F. (2002). **The Theory of Open Quantum Systems**. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199213900.001.0001>
- Bruzewicz, C. D., Chiaverini, J., McConnell, R., & Sage, J. M. (2019). Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2), 021314. <https://doi.org/10.1063/1.5088164>
- Burkard, G., Ladd, T. D., Pan, A., Nichol, J. M., & Petta, J. R. (2021). Semiconductor spin qubits. *Reviews of Modern Physics*, 93(2), 025005. <https://doi.org/10.1103/RevModPhys.93.025005>
- Caldeira, A. O., & Leggett, A. J. (1983). Quantum tunnelling in a dissipative system. *Annals of Physics*, 149(2), 374–456. [https://doi.org/10.1016/0003-4916\(83\)90202-6](https://doi.org/10.1016/0003-4916(83)90202-6)
- Doherty, M. W., Manson, N. B., Delaney, P., Jelezko, F., Wrachtrup, J., & Hollenberg, L. C. L. (2013). The nitrogen-vacancy colour centre in diamond. *Physics Reports*, 528(1), 1–45. <https://doi.org/10.1016/j.physrep.2013.02.001>
- Goold, J., Huber, M., Riera, A., del Rio, L., & Skrzypczyk, P. (2016). The role of quantum information in thermodynamics—a topical review. *Journal of Physics A: Mathematical and Theoretical*, 49(14), 143001. <https://doi.org/10.1088/1751-8113/49/14/143001>
- Google Quantum AI. (2023). Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7907), 676–681. <https://doi.org/10.1038/s41586-022-05434-1>
- Guan, X., Wang, D. M., Xie, Y. J., Fan, J., & Ji, Z. Q. (2025). Development of a cryogen-free dilution refrigerator with a cooling power of 2000 μ W at around 100 mK. *Review of Scientific Instruments*, 96(12), 123901. <https://doi.org/10.1063/5.0294894>
- Krinner, S., Storz, S., Kurpiers, P., Magnard, P., Heinsoo, J., Keller, R., Lütolf, J., Eichler, C., & Wallraff, A. (2019). Engineering cryogenic setups for 100-qubit scale superconducting circuit systems. *EPJ Quantum Technology*, 6(1), 2. <https://doi.org/10.1140/epjqt/s40507-019-0072-0>

- Müller, C., Cole, J. H., & Lisenfeld, J. (2019). Towards understanding two-level-systems in amorphous solids—a review. *Reports on Progress in Physics*, 82(12), 124501. <https://doi.org/10.1088/1361-6633/ab3a7e>
- Pauka, S. J., Das, K., Kalra, R., Vahapoglu, E., Serrano, M. A., Liles, J. D., ... Dzurak, A. S. (2021). A cryogenic CMOS chip for generating control signals for multiple qubits. *Nature Electronics*, 4*, 64–70. <https://doi.org/10.1038/s41928-020-00528-y>
- Pooley, R. A. (2005). Fundamental physics of MR imaging. *Radiographics*, 25(4), 1087–1099. <https://doi.org/10.1148/rg.254045117>
- Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
- Schlosshauer, M. (2007). **Decoherence and the Quantum-to-Classical Transition**. Springer. <https://doi.org/10.1007/978-3-540-35775-9>
- Sivak, V. V., et al. (2023). Real-time quantum error correction beyond break-even. *Nature*, 616(7955), 50–55. <https://doi.org/10.1038/s41586-023-05782-6>
- Swartz, E. T., & Pohl, R. O. (1989). Thermal boundary resistance. *Reviews of Modern Physics*, 61(3), 605–668. <https://doi.org/10.1103/RevModPhys.61.605>
- Vepsäläinen, A. P., et al. (2020). Impact of ionizing radiation on superconducting qubit coherence. *Nature*, 584(7822), 551–556. <https://doi.org/10.1038/s41586-020-2619-8>