

Microwaves in Quantum Computing

JOSEPH C. BARDIN^{1,2} (Senior Member, IEEE), DANIEL H. SLICHTER³ (Senior Member, IEEE),
AND DAVID J. REILLY^{4,5} (Member, IEEE)

(Invited Paper)

¹Department of Electrical and Computer Engineering, University of Massachusetts Amherst, Amherst, MA 01003 USA

²Google LLC, Goleta, CA 93117 USA

³Time and Frequency Division, National Institute of Standards and Technology, Boulder, CO 80305 USA

⁴Microsoft Inc., Microsoft Quantum Sydney, The University of Sydney, Sydney, NSW 2050, Australia

⁵ARC Centre of Excellence for Engineered Quantum Systems (EQUS), School of Physics, The University of Sydney, Sydney, NSW 2050, Australia

CORRESPONDING AUTHOR: Joseph C. Bardin (e-mail: jbardin@engin.umass.edu).

This work was supported in part by the National Science Foundation under Grant #1809114, and in part by the NIST Quantum Information Program.

ABSTRACT Quantum information processing systems rely on a broad range of microwave technologies and have spurred development of microwave devices and methods in new operating regimes. Here we review the use of microwave signals and systems in quantum computing, with specific reference to three leading quantum computing platforms: trapped atomic ion qubits, spin qubits in semiconductors, and superconducting qubits. We highlight some key results and progress in quantum computing achieved through the use of microwave systems, and discuss how quantum computing applications have pushed the frontiers of microwave technology in some areas. We also describe open microwave engineering challenges for the construction of large-scale, fault-tolerant quantum computers.

INDEX TERMS Semiconductor spin qubit, superconducting qubit, trapped ion qubit, quantum computing, qubit control, qubit readout, quantum-classical interface.

I. INTRODUCTION

Quantum computing and modern microwave engineering share a common ancestor in the pioneering work that led to the development of radar and related technologies in the 1940's [1]. Indeed, many of the fundamental mechanisms underpinning the generation, transmission, absorption, and detection of microwave energy were understood at that time to be governed by quantum mechanics, which describes the light-matter interaction between microwaves and atoms or molecules [2], as well as their constituent charge and spin states [3]. The leveraging of war-time radar technology and methods in the discovery of nuclear magnetic resonance in solids [4], [5] provides an ideal example of the long-standing synergy between microwave engineering and quantum systems [1].

Fundamental to this light-matter interaction is the relation $E = \hbar\omega$, which connects the angular frequency ω of microwave photons to their energy E (\hbar is the reduced Planck's constant). The quantization of microwave energy describes how microwave photons can interact resonantly with other

quantum objects, such as the quantum two-level systems (called *qubits*) that form the building blocks of quantum computers. Today, microwave technology is ubiquitous across many different quantum platforms, enabling the precise control and readout of quantum states.

Here, we review the use of microwaves in quantum computing. For the sake of brevity, rather than surveying the complete quantum computing landscape, we focus on three leading qubit technologies: trapped ion qubits, semiconductor spin qubits, and superconducting circuit qubits. We chose these technologies not only because they span the representative qubit technologies and are currently considered among the most promising of all qubit types, but also because they are heavily reliant on microwave technologies (see Fig. 1). The paper is organized as follows:

II) An overview of qubits and quantum processors is presented; the analogy between a qubit and a microwave resonator is explained, with implications discussed; and physical qubit realizations are described.




FIGURE 1. Microwave electronics for operating a 50-qubit Google quantum processor. The system generates and receives signals in the 4–8 GHz band, and was used to execute a demonstration quantum computing algorithm designed to be too complex for simulation by even the largest classical supercomputers [7]. Four racks of microwave electronics are required to control and measure the quantum processor. Abbreviations: BB—baseband, AWG—arbitrary waveform generator, SSB—single sideband. Photo credit: R. Ceselin.

- III) The coupling of a microwave source to various qubit technologies is described and typical signal and noise levels are compared and contrasted.
- IV) The quantum gate abstraction is explained, microwave techniques for implementing single and two qubit quantum gates are described, and typical hardware configurations are presented.
- V) Microwave techniques for measuring the state of a qubit are described.
- VI) Additional microwave techniques required for the operation of trapped-ion qubits are described.
- VII) Microwave innovations in quantum-limited amplifiers and non-reciprocal devices, driven by quantum computing research, are presented.
- VIII) Outstanding challenges related to microwave engineering that must be overcome to realize the full potential of quantum computing are described.

II. QUBITS AND QUANTUM COMPUTING

The fundamental information carriers in a quantum computer are quantum bits, or qubits, in analogy to the logical bits used in a classical (non-quantum) computer. Here we describe the basic characteristics and behavior of qubits, before discussing particular physical implementations of qubits and their connections with microwave technology. The interested reader can find a much more detailed exposition of this topic, written for microwave engineers, in Ref. [6].

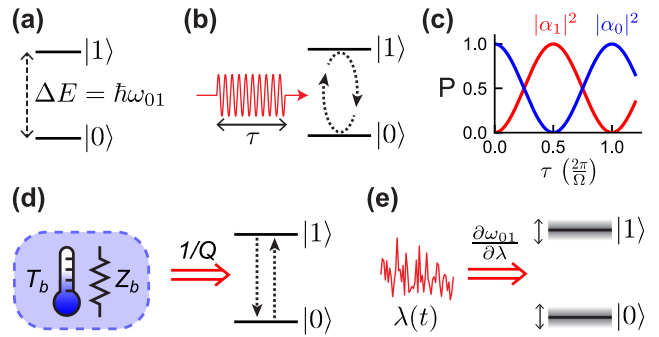


FIGURE 2. Qubit basics. (a) A qubit is composed of two quantum states $|0\rangle$ and $|1\rangle$ with an energy difference ΔE . (b) When excited resonantly at $\omega_{01} = \Delta E/\hbar$, the qubit state can be driven between $|0\rangle$ and $|1\rangle$, including into linear combinations of $|0\rangle$ and $|1\rangle$. (c) Under cw resonant driving, the qubit state undergoes so-called Rabi oscillations, where the probabilities $|\alpha_0|^2$ and $|\alpha_1|^2$ evolve sinusoidally in time. (d) Coupling to a source of dissipation at frequency ω_{01} (noise temperature T_b , impedance Z_b) with a coupling quality factor Q causes qubit state transitions on the timescale T_1 . If $k_B T_b \ll \hbar\omega_{01}$, these transitions will be dominated by the $|1\rangle \rightarrow |0\rangle$ process. (e) If the qubit frequency can be shifted by an environmental parameter λ (e.g. magnetic field), fluctuations in λ cause fluctuations in ω_{01} , dephasing the qubit state on the timescale T_ϕ .

A. QUBIT BASICS

A qubit is a quantum mechanical system with two energy eigenstates,¹ which we label as $|0\rangle$ and $|1\rangle$, with corresponding energy eigenvalues E_0 and E_1 , taking $E_0 < E_1$ (see Fig. 2(a)). The “textbook” qubit is a spin-1/2 particle (such as an electron) in a magnetic field, but as we shall see, there are many other possible ways to realize a qubit.

The $|0\rangle$ and $|1\rangle$ states are used as computational basis states, analogous to the use of “0” and “1” in classical computing. Just as the electromagnetic field can be decomposed into a linear combination of orthogonal modes (such as plane waves, spherical harmonics, or guided modes), the instantaneous state of a qubit $|\psi\rangle$ can be written as a linear combination of the two energy eigenstates, with complex amplitudes α_0 and α_1 :

$$|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle. \tag{1}$$

The customary normalization of $|\psi\rangle$ is that $|\alpha_0|^2 + |\alpha_1|^2 = 1$. To meet this criterion, we introduce the suggestive parameterization $\alpha_0 = \cos(\theta/2)$ and $\alpha_1 = e^{j\phi} \sin(\theta/2)$, where ϕ and θ are real and j is the imaginary unit, for reasons that will become clear later. For a single qubit, one is free to choose a global phase convention such that α_0 is purely real, as seen here, but in general both α_0 and α_1 are complex.

The quantum-mechanical nature of the state $|\psi\rangle$ means that the qubit can be in both states $|0\rangle$ and $|1\rangle$ simultaneously—a phenomenon known as *superposition*—in contrast to the behavior of classical bits, which can only be in one state at a time. When the state of the qubit is measured, however, the qubit state is said to “collapse” to just one of its eigenstates. The collapse is probabilistic, with the state being measured

¹The notation “ $| \cdot \rangle$ ” denotes a quantum state.

to be $|0\rangle$ with probability $P_{|0\rangle} = |\alpha_0|^2$, or $|1\rangle$ with probability $P_{|1\rangle} = |\alpha_1|^2$; for this reason, α_0 and α_1 are called *probability amplitudes*. Because of the collapse process, the post-measurement qubit state no longer contains information about α_0 and α_1 ; it has collapsed to either $|0\rangle$ or $|1\rangle$. The values of $P_{|0\rangle}$ and $P_{|1\rangle}$ can only be determined by many rounds of preparing the same initial qubit state $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ and measuring it, to build statistics on the measurement collapse probabilities $|\alpha_0|^2$ and $|\alpha_1|^2$. Although beyond the scope of this article, the complex phase $\phi = \arg(\alpha_1) - \arg(\alpha_0)$ can be determined using a procedure known as state tomography [8].

Quantum computers require more than just one qubit to perform useful computations. For N qubits, there are 2^N basis states of the system, from $|00\dots 00\rangle$ to $|11\dots 11\rangle$. While N classical bits can only be in one of the 2^N basis states at a given time, the phenomenon of superposition means that the quantum state $|\psi\rangle$ of N qubits can be a linear combination of any—or even all—of the 2^N basis states at the same time, with corresponding complex probability amplitudes $\{\alpha_{00\dots 00}, \alpha_{00\dots 01}, \dots, \alpha_{11\dots 11}\}$. Since these amplitudes have a physical interpretation in terms of probabilities of measurement outcomes, they have the normalization condition $\sum_{k=0}^{2^N-1} |\alpha_k|^2 = 1$, where k indexes the 2^N different bitstrings corresponding to the basis states.

While it is possible to access a range of superposition states by putting each qubit in its own independent superposition state, such an approach can only be used to reach a small fraction of the basis states of the qubit state vector $|\psi\rangle$, since most of the possible linear combinations exhibit correlations between the qubits. For example, consider the state $|\psi\rangle = \frac{1}{\sqrt{2}}|00\dots 00\rangle + \frac{1}{\sqrt{2}}|11\dots 11\rangle$. It is not possible to write this state as a product of separate, individual states of the constituent qubits; the state of each qubit is inextricably correlated, or *entangled*, with all the others. The state $|\psi\rangle$ above is a superposition state, where the process of measurement will cause a collapse to just one basis state. Let us consider what happens if we measure just one of the N qubits in this entangled state. If the measured qubit collapses to $|0\rangle$, then this causes all the other qubits to collapse to $|0\rangle$ as well, even though they were not measured directly, and the state becomes $|\psi\rangle = |00\dots 00\rangle$. Likewise, if the measured qubit happened to collapse to $|1\rangle$, then all the other qubits would collapse to $|1\rangle$, even though they were not measured directly, giving $|\psi\rangle = |11\dots 11\rangle$. This phenomenon of *entanglement* is a defining feature of quantum mechanics, and an essential ingredient for quantum computing.

The 2^N -dimensional state space of N qubits can hold exponentially more information than that of N classical bits, offering the hope of greatly increased computing power. However, the phenomenon of measurement collapse means that only N bits of information (a single bitstring representing the state of the N qubits after measurement), selected probabilistically by the measurement process, can be extracted from the quantum state at the end of an algorithm. In order to realize a speedup over classical algorithms, quantum algorithms generate interference between the 2^N complex

amplitudes $\{\alpha_{00\dots 00}, \alpha_{00\dots 01}, \dots, \alpha_{11\dots 11}\}$ to increase the likelihood of measuring certain output bitstrings (ones which yield the desired outcome of the computation). The action of an N -qubit quantum algorithm is analogous to the scattering matrix for a 2^N -port passive lossless microwave device, where each quantum basis state (or bitstring) is mapped to a port. An algorithm with speedup is equivalent to an S-matrix which transforms uniform input excitation at all ports into nonzero output at only a small subset of ports, through constructive and destructive interference, similarly to how a phased array can be used for beamforming, for example. Since the size of such a theoretical microwave device scales as 2^N , it would become impossible to realize for sufficiently large N . This highlights the potential for sufficiently large quantum computers to perform calculations which are classically intractable [7], [9].

Quantum algorithms create the desired interference between probability amplitudes by controlling the states of individual qubits, and generating entanglement between qubits, in the course of the algorithm. Experimentally, these tasks are usually carried out by microwave signals, or rely critically on microwave techniques. In some qubit technologies, the measurement process is also carried out by microwave signals. These topics will be detailed in Sections IV and V.

It is useful to characterize the performance of qubit state preparation, control, and measurement operations using a metric known as the *fidelity*. For each of these three tasks, there exist distinct methods to quantify the fidelity; a detailed discussion is beyond the scope of this paper [8]. In general, the fidelity can be thought of as characterizing how close the laboratory implementation of an operation is to its ideal theoretical representation, with a fidelity of 1 representing a perfect implementation and a fidelity of 0 indicating a complete failure. The presence of noise, drifts, dissipation, or miscalibration can give rise to errors in the implementation and thus cause the fidelity to be less than 1. The *error rate* characterizes the amount by which the fidelity of an operation is less than 1. Techniques for efficiently and accurately estimating fidelities and error rates, especially in larger quantum processors, are an active area of research.

B. QUBITS AS RESONATORS

One way to think of a qubit is as a high-quality-factor electromagnetic resonator, with a resonant frequency $\omega_{01} = (E_1 - E_0)/\hbar$ set by the energy difference between the qubit states, and a quality factor $Q \gg 1$. Here $Q = \omega_{01}/\gamma$, where γ is the decay rate of the energy in the qubit due to all sources of dissipation. Good qubits typically have $Q > 10^6$, and sometimes considerably higher.

Unlike ordinary linear resonators, qubits are extremely anharmonic (or nonlinear). As shown in Fig. 2(b), a resonant cw drive tone can be used to excite the qubit from the “ground” $|0\rangle$ state to the “excited” $|1\rangle$ state, but further excitation is not possible since there are no higher energy levels resonant with the drive. Continued application of the drive tone can thus only return the qubit to $|0\rangle$ again. The principle of superposition allows the qubit state to be driven to arbitrary linear

combinations of $|0\rangle$ and $|1\rangle$; accordingly, persistent resonant driving of a qubit causes sinusoidal oscillations of the probabilities $|\alpha_0|^2$ and $|\alpha_1|^2$ in time, as seen in Fig. 2(c). These oscillations are known as Rabi oscillations, and their angular frequency (called the Rabi frequency Ω) is proportional to the amplitude of the resonant drive.

As with classical resonators, a qubit's internal quality factor Q_i describes dissipation due to intrinsic loss mechanisms. As we describe in detail in Section IV, the introduction of external driving, interaction, and measurement ports—all necessary for quantum computing—will create additional loss channels which further damp the qubit resonance. Each loss channel can be identified with its own quality factor: Q_d for driving ports (where control signals are applied), Q_c for coupling ports (enabling interactions with other qubits), and Q_m for measurement ports. The total quality factor of the qubit is then given by the inverse sum:

$$\frac{1}{Q} = \frac{1}{Q_i} + \frac{1}{Q_d} + \frac{1}{Q_c} + \frac{1}{Q_m}. \quad (2)$$

The characteristic time scale T_1 over which a qubit initially in the $|1\rangle$ state will spontaneously transition to the $|0\rangle$ state² is given by Q/ω_{01} . A qubit's T_1 —analogous to the ring-down time for a high-Q resonator—is an important metric for measuring qubit performance, and should ideally be much longer than the duration of any algorithm using the qubit.

These loss channels can be thought of as arising from coupling to different sources of dissipation, or “baths” (for example, the real impedance of a control line), each with some effective noise temperature that is usually (but not always) close to the physical temperature (see Fig. 2(d)). The qubit will thermalize to these baths on the timescale T_1 . If $k_B T_b \ll \hbar\omega_{01}$, where T_b is the coupling-weighted average temperature of the baths, then the qubit will “reset” thermally to the $|0\rangle$ state by itself. For qubit frequencies of 5 GHz, this corresponds to bath temperatures in the $\lesssim 100$ mK range. Otherwise, the qubit will thermalize to some combination of $|0\rangle$ and $|1\rangle$, and must be actively reset before it is used in a computation.

In addition to loss or damping, there is a second type of decoherence we must consider. Even with a completely lossless resonator, the resonance frequency itself can still fluctuate randomly, causing the excitation in the resonator to lose phase coherence over time relative to a stable reference oscillator. If a qubit exhibits fluctuations in its resonance frequency ω_{01} , the phase information ϕ in the complex amplitudes α_0 and α_1 will be lost on a timescale T_ϕ (assuming infinite T_1). Real qubits (and real resonators) experience damping too, which also causes loss of phase information. We can define a characteristic total dephasing time T_2 over which phase information is lost as $1/T_2 = 1/(2T_1) + 1/T_\phi$. The loss of phase information destroys the interference between probability amplitudes α on which quantum algorithms rely, so T_2 should also be

²The direction of the transitions depends on the effective temperature T_b of the loss channel. If $k_B T_b \ll \hbar\omega_{01}$, the loss channel will cause relaxation from $|1\rangle$ to $|0\rangle$. However, if $k_B T_b \gtrsim \hbar\omega_{01}$, the loss channel will induce transitions in both directions.

much longer than the duration of any algorithm using the qubit.

Qubit frequency fluctuations occur when ω_{01} is sensitive to some external parameter λ . The parameter λ could be the local magnetic or electric field, for example (by analogy, for a voltage-controlled oscillator circuit, λ could be the tuning voltage). Small fluctuations $\delta\lambda$ then lead to qubit frequency fluctuations $\delta\omega_{01} = (\frac{\partial\omega_{01}}{\partial\lambda})\delta\lambda$, which causes dephasing (Fig. 2(e)). Experimentally, T_2 is maximized by reducing both the sensitivity $\frac{\partial\omega_{01}}{\partial\lambda}$ of the qubit frequency to noise and the amount of environmental noise $\langle\delta\lambda(t)^2\rangle$ present.

The qubit properties ω_{01} , T_1 , and T_2 can vary widely between different qubit technologies, as we will see in the next section. We will expand further on the analogy between qubits and high-quality-factor resonators in Section III.

C. PHYSICAL REALIZATION OF QUBITS

Just as classical bits can have many different physical realizations—the voltage on the gate of a transistor, the spin orientation of a small magnetic domain on a hard disk, the reflectivity of a small region of an optical storage medium—qubits can have different physical implementations as well. In this paper, we focus on three leading physical implementations of qubits: trapped atomic ions, spins in semiconductors, and superconducting circuits. Below, we briefly explain the fundamentals and properties of these different types of qubits, all of which can have ω_{01} in the microwave region of the spectrum, as seen in Fig. 3.

1) TRAPPED ION QUBITS

Qubits realized in the quantum states of atomic ions trapped in ultra-high vacuum are one of the most mature and high-fidelity quantum technologies [10]–[12]. Because ions have a net charge, they can be readily trapped and held in isolation in vacuum using electromagnetic fields. For quantum computing applications, the ions are usually trapped using so-called linear rf Paul traps [11], which confine charged particles—for hours to months, depending on parameters—using a combination of static electric fields and oscillating radio-frequency (typically between 20 MHz and 150 MHz) electric fields. These fields are generated by applying dc and/or rf potentials to sets of trapping electrodes; some example ion traps are shown in Fig. 4. The largest traps have centimeter-scale electrodes made in a machine shop, usually held together with insulating ceramic parts in a three-dimensional geometry (Fig. 4(a)). Intermediate-scale three-dimensional traps, with electrode dimensions down to hundreds of microns, can be made by depositing patterned metal films on laser-cut or etched insulating substrates (Fig. 4(b)). Two-dimensional traps, known as surface-electrode ion traps [13], are made on planar substrates using microfabrication techniques, and have typical electrode dimensions from ~ 100 μm down to a few μm (Fig. 4(c)). Numerous groups, including commercial quantum computing entities, are pursuing surface-electrode




FIGURE 3. Relevant ranges of the electromagnetic spectrum for quantum computing. We indicate the regions of the spectrum typically used by different qubit technologies for control and readout. The carrier frequency for trapped ions can be as low as a few MHz. The narrow trapped ion carrier frequency bands at ~ 10 GHz and ~ 12.6 GHz correspond to $^{133}\text{Ba}^+$ and $^{171}\text{Yb}^+$, respectively. Signals at the carrier frequencies are modulated with instantaneous IF/baseband bandwidths shown in the right panel. These technologies also all use direct baseband control signals with bandwidths as shown in the right panel.




FIGURE 4. Trapped ion qubits. (a) Machined 3D trap, (b) laser-machined 3D trap, and (c) microfabricated surface-electrode trap with integrated microwave antenna structures for qubit control (blue). Ions (shown schematically as red spheres in (c), each in a separate trapping potential well) are held at reconfigurable positions along the red dotted lines, between ((a) and (b)) or above (c) the trap electrodes. (d) Camera image of chain of trapped ions, fluorescing from laser excitation. Ion-ion spacing is a few μm . (e) Energy level diagram of generic trapped ion Zeeman qubit. (f) Energy level diagram of $^9\text{Be}^+$ hyperfine qubit with $B_0 > 0$ (the levels are degenerate at $B_0 = 0$, shown as grey dotted lines), and one particular choice of qubit levels. In $^9\text{Be}^+$, $\omega_{\text{HFS}}/2\pi = 1.25$ GHz. In (e) and (f), the optical transitions to excited electronic states (blue arrows, not to scale) are used for laser cooling, qubit state preparation, and qubit readout. Photo credits: Ion Storage Group/NIST.

traps as a path toward large-scale trapped-ion quantum computing, due to the ability to make complex trap designs with many different trapping regions to hold large numbers of ions [14], [15].

Ions are typically loaded into traps by electron impact ionization or resonant photoionization of a flux of neutral atoms from an in-vacuum thermal oven or laser ablation target [12]. When more than one ion is held in a trap, the mutual Coulomb repulsion between ions keeps them spatially separated. When laser-cooled [10], [16], the ions form a static Coulomb crystal with ion-ion spacings on the order of a few μm , as seen

in Fig. 4(d). The motion of the ions in this configuration is strongly coupled by the Coulomb force, and is described as a set of collective normal modes of ion motion, whose resonant frequencies are typically between 500 kHz and 10 MHz. Traps can have multiple spatially separated potential wells, each containing such a Coulomb crystal.

Trapped ions used in quantum computing applications typically have a single valence electron, with all other electrons in closed shells. Among the ion species most commonly used for quantum computing applications are $^9\text{Be}^+$, $^{25}\text{Mg}^+$, $^{40}\text{Ca}^+$, $^{43}\text{Ca}^+$, $^{88}\text{Sr}^+$, $^{133}\text{Ba}^+$, $^{138}\text{Ba}^+$, and $^{171}\text{Yb}^+$, although numerous others have been employed as well. Below, we describe several possible choices of qubit levels from among the many quantum states of these ions.

When the ion species has no nuclear spin, a qubit can be realized using the two spin states of the ground state valence electron in the presence of an external magnetic field. This type of trapped ion qubit is known as a Zeeman qubit, and is shown in Fig. 4(e). The qubit resonance frequency is proportional to the external magnetic field B_0 according to $\omega_{01}/2\pi = (\gamma_e/2\pi)|B_0|$, where $\gamma_e/2\pi \approx 28$ GHz/T is the electron gyromagnetic ratio. Trapped ion Zeeman qubits are typically operated in magnetic fields of less than 1 mT, giving qubit frequencies of ~ 10 MHz. For ion species with nonzero nuclear spin, such as $^9\text{Be}^+$ or $^{171}\text{Yb}^+$, the hyperfine interaction between the nuclear spin and the valence electron spin gives rise to two manifolds of hyperfine states in the ground electronic state, as shown in Fig. 4(f). At low magnetic fields ($B_0 \lesssim 50$ mT), these manifolds are separated by the hyperfine splitting $\omega_{\text{HFS}}/2\pi$, which ranges from 1.25 GHz for $^9\text{Be}^+$ to 12.6 GHz for $^{171}\text{Yb}^+$, with higher-mass ions having larger splittings. A small magnetic field B_0 (typically < 1 mT, but sometimes up to \sim tens of mT) is applied such that each hyperfine state has a unique energy. A trapped ion hyperfine qubit consists of two such states, usually chosen to be in separate hyperfine manifolds. The qubit resonance frequency ω_{01} may differ by up to several hundred MHz from ω_{HFS} for the range of magnetic fields listed above, depending on the choice of qubit states. Zeeman and hyperfine qubits can be manipulated directly using rf or microwave magnetic fields, either launched in free space by distant antennas or horns, or

from local antenna structures fabricated in the trap (as seen in Fig. 4(c)).

Trapped ions in ultra-high vacuum are isolated from the nearest surfaces and bulk materials by tens to hundreds of μm . As a result, electric and magnetic field noise at the ion are orders of magnitude lower than typically seen inside or on the surface of solids; combined with the relatively weak coupling of the quantum states of the ion to external fields, this means that trapped ion qubits do not thermalize to the environment rapidly. Typical T_1 values for Zeeman and hyperfine qubits are years. As a result, optical pumping is used to initialize the internal states of the ions, and laser cooling is used to bring the ion motion near its quantum mechanical ground state [10], [12]. This very slow thermalization also means that the temperature of the trap electrodes and the vacuum chamber need not satisfy $k_B T \ll \hbar\omega_{01}$ to achieve quantum behavior. However, cryogenic operation of ion traps (in the 4 K to 10 K range) can be useful for increasing ion lifetime in the trap by cryopumping background gas, and for reducing electric field noise that heats and decoheres the ion motion [12].

Trapped ion qubits can be dephased by magnetic field fluctuations. Zeeman qubits are directly sensitive to magnetic field fluctuations, with $T_2 \sim$ tens of ms, but T_2 values of up to 300 ms have been achieved with appropriate magnetic field shielding [17]. For hyperfine qubits, it is possible to choose B_0 such that a particular hyperfine transition is insensitive to magnetic field noise to first order. Such a qubit is known as a “clock” qubit, so named because field-insensitive transitions generally have very long dephasing times and are therefore ideal for realizing microwave-frequency atomic clocks. Bare clock qubit T_2 values are usually $\gtrsim 1$ s, but values as high as 50 s have been reported [18]; performing a type of qubit “chopping” (called *dynamical decoupling*) to counteract $1/f$ magnetic field noise can yield T_2 in excess of an hour [19]. Measurements of qubit coherence on these timescales are generally limited by the frequency stability and drift of the microwave reference oscillator to which the qubit is compared [19]–[21].

The interested reader is referred to Ref. [12] for further details on trapped ion quantum computing.

2) SEMICONDUCTOR SPIN QUBITS

The spin degree of freedom in solids provides another potential platform for scalable quantum computing systems. Nuclear spins in silicon, for instance, can exhibit hours-long T_1 times [22]. In contrast to the vacuum of a trapped ion, spin qubits are embedded in solids and surrounded by other atoms, many of which may interact with the spin qubit in uncontrolled ways [3]. Fortunately, these interactions are relatively weak in materials that contain few nuclear spins such as silicon [23], silicon-germanium [24], and carbon materials such as diamond [25].

Highly isolated qubits that are only weakly interacting with their environment are also, in general, weakly coupled to any means of control. Weak coupling to control fields results in

slow quantum gate times, potentially cancelling any advantage afforded by the long coherence times. For this reason many different ‘flavors’ of spin qubit have been devised, with trade-offs between controllability, device complexity, and sensitivity to charge or voltage noise. In this review, we will limit our discussion to spin qubits based on confined electrons or holes in semiconductors.

Modern nanofabrication makes it possible to confine and detect single electron spins in ‘zero-dimensional’ nanostructures referred to as quantum dots (QDs) [26]. The potential that confines the electron (or hole) is produced electrostatically via gate electrodes on the surface of a semiconductor, enabling the number of electrons on a dot and their coupling to the neighboring dots and reservoirs to be tuned by varying gate voltages. The ability to confine, manipulate, and detect single spin states on quantum dots is largely a consequence of the Coulomb blockade of charge, an electrostatic phenomenon arising when the energy to charge a capacitor C by a single electron charge e , $E = e^2/2C$, is larger than thermal energy $k_B T$. For sub-micron devices with self-capacitance in the attofarad range, the energy scale for Coulomb blockade requires temperatures below a few kelvin. This necessitates the use of dilution refrigerators for operating spin qubit systems, although work is underway to operate at elevated temperatures [27], [28].

The initial proposal by Loss and DiVincenzo [29] for a spin-based quantum computer assumed arrays of coupled quantum dots, each hosting a single electron spin. A large external magnetic field B_0 then sets the energy difference between the two spin states aligned or anti-aligned with B_0 , much as with a trapped ion Zeeman qubit. The qubit resonance frequency is given by $\omega_{01}/2\pi = (\gamma_e/2\pi)|B_0|$. Single electron spin qubits are typically operated in tesla-scale magnetic fields in order to ensure $\hbar\omega_{01} \gg k_B T$, with $\omega_{01}/2\pi \sim 1 - 50$ GHz.

Kane proposed exploiting the exceedingly long coherence of phosphorous donors in isotopically purified ^{28}Si by coupling their nuclei to localized electron spins for single qubit addressing, two-qubit coupling, and readout [22]. Any such qubit architecture requires methods for the precise placement of single atomic donors in a solid. Despite significant experimental progress since Kane’s original proposal [30], the MHz (rather than GHz) resonance frequencies of nuclei present a major challenge, leading to kHz clock rates for a quantum computer. The potential for nuclear spins to be used as quantum memories, however, appears more promising [30].

The challenge of requiring GHz-frequency magnetic fields for single spin manipulation can be overcome at the expense of requiring a two-electron system for a single qubit. Here, double quantum dots are used to host two tunnel-coupled electrons, and the qubit is created by the energy splitting of the spin singlet state $|S\rangle$ and one of the three spin triplet³ configurations ($|T_+\rangle$, $|T_0\rangle$, $|T_-\rangle$) [31]. Coupling singlet-triplet

³For two spins, each with basis states $|\uparrow\rangle$ and $|\downarrow\rangle$, these states are $|S\rangle = \frac{1}{\sqrt{2}}(|\downarrow\uparrow\rangle - |\uparrow\downarrow\rangle)$, $|T_+\rangle = |\uparrow\uparrow\rangle$, $|T_0\rangle = \frac{1}{\sqrt{2}}(|\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle)$, $|T_-\rangle = |\downarrow\downarrow\rangle$.




FIGURE 5. A five quantum dot device in GaAs with additional quantum dot charge sensors either side. White arrows indicate microwave currents used in readout. Blue and orange ovals are electrons. The pairs of blue ovals represent two S-T qubits.

(S-T) spin qubits is challenging, and work is underway to devise various architectures to facilitate two-qubit gates. A pair of S-T qubits, each requiring a double quantum dot, is shown in Fig. 5. A multi-electron dot (shown in orange) provides a means of coupling the qubits [32].

One can extend the idea of using multiple spins to define a qubit by implementing three exchange-coupled electrons. This approach creates a qubit via the relative orientation of spins, rather than the alignment of spins to an external magnetic field. Requiring three quantum dots as host platform, the exchange-only (E-O) qubit [33] (and its ac variant, the resonant exchange qubit [34]) can be controlled entirely with baseband gate voltages rather than microwave magnetic fields.

The coherence times of semiconductor spin qubits are $T_1 \sim 0.1\text{--}1$ ms and T_2 of tens to hundreds of μs for all of the variants described above. For a detailed review of spin qubit technology, we refer the reader to [35].

3) SUPERCONDUCTING QUBITS

Unlike the qubit variants described above, whose degrees of freedom are those of single electrons and atomic nuclei, superconducting qubits are macroscopic devices that are defined at the circuit level and implemented using nominally lossless capacitors, inductors, and Josephson junctions (JJs). When operated at a low enough physical temperature—typically in the low tens of millikelvins—these circuits display coherent quantum mechanical behavior [36], as necessary for use in a quantum processor. Since they are constructed using lumped and/or distributed circuit elements, the properties of superconducting qubits can be engineered similarly to classical circuit structures. As such, a rich family of quantum devices can be realized in this technology platform. Owing to the engineerable nature of these monolithically fabricated devices, the field of superconducting qubit technology has attracted considerable attention since the first device was realized by




FIGURE 6. The transmon qubit. (a) Schematic diagram. The “X” symbol represents a Josephson junction. (b) Energy diagram. The unequal energy spacings are due to the qubit nonlinearity.

Nakamura *et al.* a little over two decades ago [37]. Here, we focus on one particular type of superconducting qubit, the transmon [38], currently used in commercial quantum computing efforts [7]. For a detailed review of superconducting qubit technology, we refer the reader to [39].

A transmon qubit is a nonlinear microwave LC resonator, constructed by shunting a JJ with a capacitance C_Q , as shown schematically in Fig. 6(a). The nonlinearity arises from the JJ, which behaves as a current-dependent inductance $L_J = L_{J0}/\sqrt{1 - I_J^2/I_C^2}$, where $L_{J0} = \Phi_0/2\pi I_C$ is the zero-bias inductance of the JJ, $\Phi_0 = \pi\hbar/e$ is the magnetic flux quantum, e is the elementary charge, I_J is the current through the JJ, and I_C is the critical current of the JJ. The value of I_C is determined by the JJ geometry and is typically about 40 nA for transmons, corresponding to $L_{J0} \approx 8$ nH [40].

The nonlinearity causes the transmon’s resonant frequency to decrease in proportion to the energy stored in the resonator, leading to an energy level diagram of the form shown in Fig. 6(b). The non-uniform nature of the energy spacings is referred to as *anharmonicity* and, while quantum engineers typically refer to the transmon qubit as weakly anharmonic, it is actually very non-linear in comparison to typical microwave components. The resonant frequency shifts by an amount $\eta = -e^2/2\hbar C_Q$ for each microwave photon added to the energy in the transmon, such that all transitions between neighboring transmon energy levels are at different frequencies. For typical component values, $\omega_{01}/2\pi$ and $|\eta/2\pi|$ are in the range of 4–8 GHz and 150–300 MHz, respectively. To take an example, for $\omega_{01}/2\pi = 6$ GHz and $|\eta/2\pi| = 200$ MHz, the addition of a single microwave photon of energy 4×10^{-24} J causes a drop in the transmon’s resonant frequency of over 3%.

Transitioning from energy E_i to energy E_j and vice-versa requires coupling energy into the transmon at frequency ω_{ij} , so one can treat the transmon as an ideal two-level qubit if caution is taken never to excite it in a way which results in leaving the ($|0\rangle, |1\rangle$) manifold—that is, one must not drive the device at ω_{12} . Notably, the transmon anharmonicity can be engineered through the choice of C_Q (while selecting the appropriate JJ sizing to obtain a desired ω_{01}). However, smaller C_Q results in an increased sensitivity to $1/f$ charge noise, resulting in a practical upper limit on $|\eta|$. As we will see, this places a constraint on the spectral content of the signals used to drive the ω_{01} transition.

The relaxation time constant T_1 of a transmon is limited by materials losses; recent work has demonstrated values of

T_1 as high as 300 μs for isolated planar qubits fabricated on a sapphire substrate [41]. Unfortunately, qubits used in a quantum processor tend to be coupled to additional loss channels, and the best values of T_1 reported for a ≥ 25 qubit quantum processor are about a factor of three lower [42]. These losses can occur through couplings to local defects that behave as two level systems (TLSs), among other mechanisms. The transition frequencies of such TLSs can be time dependent, causing T_1 to vary with time [43]. Mitigating such effects is an active area of research.

The dephasing time constant T_ϕ of single-JJ transmons is typically limited in part by fluctuations in I_C , resulting in time-dependent variations in L_{J0} . Transmons used in some of today's $\gtrsim 25$ qubit quantum processors have T_2 as high as 100 μs [42]. However, many contemporary quantum processor architectures leverage frequency tunable transmons (described below), and these devices have considerably lower T_2 , due to coupling to magnetic flux noise.

III. INTERFACING A MICROWAVE SOURCE TO A QUBIT

Transitioning between the $|0\rangle$ and $|1\rangle$ states of trapped ion, single-electron spin, and transmon qubits requires exciting the qubit on resonance, which means that we need some mechanism to couple microwave energy to the device. This can be done either electrically or magnetically, and since a qubit can be thought of as a microwave resonator, we can also think of coupling to it just as one would couple to any other microwave resonator.

In the next sections, we describe some of the considerations related to microwave drive and deterministic state control of qubits in a quantum processor.

A. DRIVE COUPLING QUALITY FACTOR, Q_d

The qubit-drive coupling can be quantified in terms of a drive coupling quality factor Q_d , defined as the contribution to the loaded quality factor of the qubit resonance due to dissipation in the impedance of the drive source. The value of Q_d sets an upper limit on the qubit's relaxation time constant,

$$T_1 \leq \frac{Q_d}{\omega_01} = T_{1,d}. \quad (3)$$

Qubits are very under-coupled to the drive source so that $T_{1,d} \gg T_1$. However, the degree to which the drive is under-coupled to the qubit varies drastically from technology to technology.

In superconducting qubits, where the qubit is engineered into a circuit environment, the coupling quality factor associated with the microwave drive can be engineered just like in any passive planar circuit. For example, with the capacitive coupling of Fig. 7(a), $Q_d \approx C_Q/(C_D^2 Z_0 \omega_{01})$; thus, Q_d can be set by choosing the coupling capacitor C_D . Materials properties limit the internal quality factor Q_i of today's state-of-the-art transmons to about 4×10^6 [41], so Q_d is typically designed to be about an order of magnitude larger such that it does not limit the qubit T_1 . For the typical case of capacitive coupling (Fig. 7(a)), $Q_d \approx 4 \times 10^7$ corresponds to a




FIGURE 7. Engineering control and measurement ports into the transmon circuit. (a) An XY control port that is capacitively coupled to the qubit permits microwave drive of the $|0\rangle \leftrightarrow |1\rangle$ transition. (b) Replacing the single JJ with a flux-biased SQUID permits control of ω_{01} via a current bias. (c) A linear resonator is coupled to the qubit for dispersive qubit state readout. (d) Photograph showing each of these circuit techniques being simultaneously employed. Capacitances C_D and C_g are those between the coplanar waveguide lines marked “XY control” and “readout resonator” and the qubit, respectively. In addition, the inductive coupling between the Z control line and the qubit is visible in the expanded view.

coupling capacitance of 30 aF [40]. Fortunately, this coupling capacitor—which limits the T_1 of a 6 GHz qubit to just over 1 ms—can be readily designed using modern electromagnetic design tools.

The coupling of microwave fields to semiconducting spin qubits, and to trapped ion hyperfine and Zeeman qubits, can be thought of as an inductive coupling: the microwave magnetic field couples to the electron spin, whose magnetic moment is fixed by nature and is “atom-sized”. By contrast, superconducting qubits use an electric field coupling to the qubit circuit, where the effective electric dipole moment can be engineered to be much larger than “atom-sized” by increasing the dimensions of the qubit circuit. As an analogy, one can think of the drive as coupling either to an extremely small loop antenna or to a large dipole antenna.

This difference means that if one puts the different types of qubits at the same distance from a propagating electromagnetic wave on a drive line, the Q_d would be roughly 10^8 times higher for semiconductor spin qubits and trapped ion hyperfine and Zeeman qubits than for superconducting qubits, given typical superconducting qubit parameters. The Q_d can also be made larger or smaller by increasing or decreasing (respectively) the distance between the qubit and the drive line, because of the fall-off of drive field strengths.

For semiconductor spin qubits, values of Q_d in the 10^{13} to 10^{15} range have been reported [44], [45]; these qubits were located within ~ 100 nm of the driving transmission line. For

trapped ions, the smallest reported experimental values of Q_d are $\sim 10^{19}$ [46], using surface-electrode traps with integrated near-field antenna structures where the ions are $\sim 30 \mu\text{m}$ from the antenna (see Fig. 4(c)). Many experiments use large three-dimensional traps with microwave horns outside the vacuum chamber, giving substantially weaker coupling and Q_d values up to $\sim 10^{24}$ [19].

B. RABI OSCILLATION FREQUENCY

Once the value of Q_d is known, we can readily determine the Rabi oscillation frequency as a function of average available power during the pulse P_{av} , referenced to the qubit drive port

$$\Omega = 2\sqrt{\frac{P_{\text{av}}}{\hbar Q_d}}. \quad (4)$$

The Rabi frequency is an important metric, in that π/Ω is the time required to flip between the $|0\rangle$ and $|1\rangle$ states. From (4) we see that the required available power referenced to the drive terminal to achieve a given Rabi frequency is directly proportional to the coupling quality factor. It is also important to emphasize that the Rabi frequency is proportional to signal amplitude (as opposed to power) since the qubit is a coherent device.

Since the qubit to drive line coupling is strongest for superconducting qubits, the available power required to achieve a given Rabi frequency is lowest. For instance, achieving a Rabi frequency of 50 MHz with a superconducting qubit having Q_d of 4×10^7 requires an available power of about -70 dBm. Compare this to a trapped ion qubit with $Q_d \approx 10^{19}$, where roughly $+40$ dBm is required to achieve Rabi frequencies in this range (the highest reported Rabi frequency for a trapped ion qubit using microwaves is 26 MHz [46]). Most microwave single-qubit gates for trapped ions are implemented with Rabi frequencies below ~ 100 kHz. For semiconductor spin qubits, achieving a Rabi frequency of 50 MHz requires an available power in the range of -16 dBm to $+4$ dBm [44], [45], referenced to the drive line, with the exact level depending upon the mode of inductive coupling to the qubit. This increased sensitivity compared to trapped ions is simply due to the closer proximity of the feed structure to the qubit.

C. EFFECT OF NOISE COUPLED THROUGH MICROWAVE DRIVE LINE

One must also consider the effect of thermal noise coupled to the qubit through the microwave drive line. Since this port is used to drive transitions between $|0\rangle$ and $|1\rangle$, noise at ω_{01} injected to the qubit through this channel will also induce transitions, leading to a transition rate [47]:

$$R_{\uparrow\downarrow,d} = \frac{S_{\text{av}}}{\hbar Q_d}, \quad (5)$$

where S_{av} is the spectral density of the noise power at the qubit frequency that is available at the qubit drive port. For a fixed value of S_{av} , the transition rate is inversely proportional to the coupling quality factor.

In general, we do not want noise on the drive line to limit coherence, so it is useful to put the required noise levels in context for each of the technologies under consideration. To begin, let us consider the limit in which relaxation through the drive line produces the same amount of decoherence as noise emitted from the drive line—that is, $R_{\uparrow\downarrow,d} = 1/T_{1,d}$. Remarkably, keeping the rate of decoherence for these two mechanisms equal requires the effective noise temperature presented by the drive line ($T_{e,d}$) be kept at the single photon level ($T_{e,d} = \hbar\omega_{01}/k_B$), independent of the coupling strength from the drive line to the qubit.

For the case of superconducting qubits, the rate of relaxation through the drive line is typically within an order of magnitude of the qubit's intrinsic relaxation rate, and it is essential that noise on the drive line not further decohere the qubit. Therefore, the limit described above is relevant, which means that it is essential to attenuate the thermal noise floor well below the single-photon noise temperature (≈ 300 mK at 6 GHz). As such, the microwave control lines typically feature heavy attenuation, with the final 20–30 dB of loss thermalized to the lowest temperature stage.

Since semiconductor spin qubits are coupled less strongly to the microwave drive source in comparison to superconducting qubits, they are also considerably less sensitive to noise. For the range of Q_d described in Section III.III-A, the value of $T_{1,d}$ for $\omega_{01}/2\pi = 5$ GHz is between 5 minutes and 9 hours. Since the coherence of semiconductor spin qubits is limited by other mechanisms to timescales considerably shorter than this, the single photon limit described above is not applicable. Limiting the rate of qubit transitions due to noise on the microwave drive line to one per second requires limiting the effective noise temperature of the drive line to between 75 K and 7,500 K, depending upon the qubit coupling to the drive line.

Trapped ion qubits are so weakly coupled to the microwave drive that a considerable amount of noise (in absolute terms) can be tolerated on the drive line. For instance, in a worst-case scenario using the lowest reported values of $Q_d \approx 10^{19}$ and requiring the highest reported Rabi frequency of ≈ 25 MHz, a typical effective drive line noise would be $S_{\text{av}} = -120$ dBm/Hz, or a noise temperature of $\approx 7 \times 10^7$ K. This would limit the qubit T_1 to approximately 1 second. However, since values of Q_d are typically orders of magnitude higher, and substantially lower Rabi frequencies are used, in practice the limit on T_1 due to the drive line noise would be minutes or hours.

The noise requirements are drastically different for each of the technologies, so it is worth asking whether there is any common ground between the signal and noise requirements. It turns out that it is in fact possible to relate the signal-to-noise ratio on the drive line to the Rabi frequency and transition rates:

$$\frac{P_{\text{av}}}{S_{\text{av}}} = \frac{\Omega^2}{4R_{\uparrow\downarrow}}. \quad (6)$$




FIGURE 8. Diagram of a simple quantum algorithm. The horizontal axis represents time, positive towards the right. Each horizontal black line represents a qubit (labeled with subscripts as “a” and “b”). Each green box represents a control operation on a single qubit, while the orange box represents a control operation that entangles two qubits. The blue boxes at far right indicate measurement of the qubits at the end of the algorithm. By preparing the identical initial state, running the algorithm, measuring, and then repeating this cycle many times, one can build up statistics about $P_{|00\rangle}$, and the probabilities of the other three possible outcomes. These probabilities represent the result of the algorithm.

Thus, once a desired Rabi frequency and Rabi frequency to transition rate ratio (effectively the average number of transitions which can be coherently driven before a noise-induced transition occurs) are determined, the required signal to noise per unit bandwidth is easily calculated using this universal relationship. Shaped pulses have a peak to average ratio which is greater than unity, so additional margin is required if using such control waveforms.

IV. COHERENT CONTROL OF QUANTUM PROCESSORS USING MICROWAVE TECHNIQUES

Thus far, we have described how one might drive a single qubit between the $|0\rangle$ and $|1\rangle$ states, but to perform quantum computing, we need coherent control of the full multi-qubit complex state vector. This section describes the role of microwaves in this process.

To get a sense of the type of control needed to implement a quantum algorithm, we will begin by considering the simple quantum algorithm shown in Fig. 8, which is used to generate an entangled state known as a Bell state: $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. The quantum state can be written as a four-element vector representing the four complex amplitudes α_{00} , α_{01} , α_{10} , and α_{11} . Control operations, called gates, can be represented by matrices that act on the state vector, as shown in Fig. 8.

The algorithm begins by resetting both qubits (labeled “a” and “b”) to the $|0\rangle$ state. Then a series of quantum gate operations are carried out. First, each qubit is placed in a superposition state by applying a so-called Hadamard gate H to each of the qubits. After this step, the two qubits are in an equal superposition of each of the four possible basis states. Next, the qubits are entangled via a controlled Z (CZ) gate, which inverts the sign of α_{11} while leaving the other amplitudes unaltered. While it is not obvious from the measurement statistics (which are unaltered by the application of the CZ gate), this is an entangled state, since it is no longer possible to describe the joint state $|\psi_{ab}\rangle$ as a product of single qubit states.

A Hadamard gate is then applied to the second qubit, which produces constructive and destructive interference between α_{00} and α_{01} , as well as α_{11} and α_{10} , resulting in the production of the desired entangled Bell state. Finally, a measurement is carried out on both qubits.

While the algorithm described above only involves a pair of qubits, it turns out that a library of gate operations giving full state control of a single qubit (described by 2×2 unitary matrices), combined with a single two-qubit entangling gate (described by a 4×4 unitary matrix), is sufficient to implement a universal quantum algorithm—that is, one can decompose any arbitrary $2^N \times 2^N$ unitary operator into a sequence of these basic operations, each of which is applied to either a single qubit or a pair of qubits [8]. This so-called *universal gate set* can be thought of in analogy to how all digital logic operations in a classical computer can be constructed from NAND gates, for example. Just as NAND gates are only one of many possible choices of a universal gate for classical computing, there are many choices for the universal gate set used in quantum computing; the particular choice of the universal gate set varies from technology to technology, since each technology has its own particularly convenient set of “native” gate operations. However, in contrast with classical computing, where a gate is thought of as a physical object implemented with transistors to which bits are brought to carry out logical operations, a quantum gate is an operation applied directly to a qubit or a pair of qubits *in situ*. Quantum gate operations are often carried out using microwave techniques.

When thinking about qubit control and measurement (to be discussed in Section V), a natural question to ask is: how good must our control and measurement be? To what extent can we tolerate errors in either? The answer depends on a variety of factors, but in general, the lower the errors, the larger the algorithm that can be run successfully, so striving for lower error is important. State-of-the-art error rates for control and measurement are currently in the range of 10^{-2}

to 10^{-6} per operation, depending on the type of operation and the physical qubit implementation. These are much higher than typical error rates in classical computing hardware. The primary reason for this difference is that analog fluctuations on a digital signal in classical computing must have a certain minimum amplitude before the signal crosses a digital logic level and gives rise to a logical error, affording high intrinsic noise immunity. By contrast, any amount of noise or miscalibration in the operations of a quantum computer can affect the continuously-variable amplitudes α of the computational basis states, potentially altering the result of the computation.⁴

Since it is never possible to eliminate errors completely, error correction techniques must be used for large-scale computations. Quantum error correction (QEC) is a large field of active research, and we will not attempt to give details in this paper, directing the reader instead to [8], [48], [49] and references therein. However, a general rule of thumb is that control errors below a threshold of roughly 10^{-4} per gate are low enough for most QEC protocols [50], and some can tolerate errors as high as roughly 1% per gate [51]. Another general rule of thumb is that the larger the gate error rate, the more resources (in terms of qubits, gate operations, and measurements) are required to implement the QEC protocol, an overhead that can become very cumbersome as error rates near the threshold. A practical target for average control error rates is 10^{-4} per gate. Since control errors are an aggregation of many different physical error mechanisms, including decoherence mechanisms intrinsic to the qubits themselves, meeting this goal means that each error contribution should be significantly lower; here we consider the goal of a maximum error contribution of 10^{-5} for each of the control error mechanisms. Readout errors should generally be at similar levels, although certain error correcting codes can tolerate readout errors as high as $\sim 1\%$.

A. SINGLE QUBIT GATES

While Rabi oscillations describe the response of a qubit to the amplitude of a resonant drive signal, a qubit is a coherent device, and just like the IQ outputs of a direct conversion receiver depend upon the relative phase relationship between the LO and RF signals, the response of a qubit also depends upon the phase of the drive signal. To help understand this relationship, we consider the single-qubit state vector, $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{j\phi} \sin(\theta/2)|1\rangle$. Since the qubit state is completely described by angles θ and ϕ , it can be thought of as a vector terminating on the surface of a unit sphere, which is referred to as the Bloch sphere (Fig. 9).⁵ In this

⁴We stress that quantum computing is not analog computing, in the sense of circuit-based analog classical computers. A quantum computer uses a discrete set of operations performed on a discrete set of basis states, and gives a digital output; control and measurement errors can be thought of as digital errors appearing with some probability, a fact which underlies the ability to perform quantum error correction [8].

⁵While the Bloch sphere representation of a single-qubit state is a valuable tool for gaining intuition into single-qubit gate operations, once two qubits are entangled, we can no longer separate their states and the Bloch sphere picture fails to be meaningful.




FIGURE 9. Bloch-sphere representation of a single-qubit state.

representation, the $|0\rangle$ and $|1\rangle$ states map to the north and south poles, whereas all other points correspond to unique superposition states. In the Bloch sphere picture, single-qubit quantum gate operations can be thought of as rotations of the qubit state vector. To construct a universal single-qubit gate set, we must be able to perform deterministic rotations of the qubit state about the X , Y , and Z axes.

1) XY GATES

Single-qubit gate operations are typically classified into two types of gates, XY and Z , which each typically have their own physical implementation. As the name suggests, XY gates produce rotations about an axis in the XY plane of the Bloch sphere. Since resonant microwave excitation of the qubit produces oscillation between the $|0\rangle$ and $|1\rangle$ states, microwave pulses can be used to mediate XY gates. As the Rabi frequency is proportional to the drive amplitude, we can rotate the state by a deterministic amount simply by controlling the integrated envelope amplitude of the microwave pulse. This can be accomplished by controlling the envelope amplitude and pulse duration. The axis of rotation is set by the microwave carrier phase, so adding control of this degree of freedom allows access to an arbitrary set of XY gates. Additionally, detuning the drive from ω_{01} causes the axis of rotation to tilt away from the XY plane, so drive frequency can be used as an additional degree of freedom. Achieving a 180° rotation (also called a ‘ π pulse’) in a duration τ_g nanoseconds while limiting the impact of each of the error mechanisms to below 10^{-5} requires control of the integrated envelope amplitude, carrier phase, and carrier frequency offset to better than 0.25% , 0.22° , and $(2/\tau_g)$ MHz, respectively.

In general, it is desirable that gate operations be carried out as quickly as possible to limit the impact of decoherence on the achieved gate fidelities. However, the spectral width of a control pulse is inversely proportional to pulse duration, so it is essential to design these pulses properly to avoid driving undesired off-resonant transitions. This is especially true when working with transmon qubits, where the typical separation of ω_{01} and ω_{12} is about $2\pi \times 200$ MHz, on the order of Ω . The leakage rate to an undesired transition ω_{ij} can be estimated as the relative energy at ω_{ij} to that at ω_{01} [52], so pulse shaping techniques are typically employed when driving transmons to

optimize the tradeoff between pulse duration and ω_{12} drive. Common envelope waveforms include Gaussian and raised cosine shapes, which have much reduced frequency-domain sidelobes in comparison to a rectangular envelope. These simple envelopes are sufficient to achieve pulse durations as short as about 20 ns, but reaching shorter gate durations with transmon qubits requires further waveform optimization. First, to further suppress drive of ω_{12} , one can employ the derivative removal by adiabatic gate (DRAG) technique, in which a notch is generated at ω_{12} by adding a quadrature derivative term to the baseband envelope [53], [54]. While this takes care of the ω_{12} leakage term, moving to shorter gate durations requires larger amplitudes and, due to the ac Stark effect, the effective value of ω_{01} becomes amplitude-dependent. As such, a time-varying detuning also must be applied to the microwave carrier signal [55].

In trapped ion qubits, XY gate fidelity is typically limited by errors in the microwave pulse parameters as described above, rather than qubit decoherence during the gate, especially for clock qubits. For this reason, high-fidelity microwave gate pulses are generally of longer duration, $\sim 1 \mu\text{s}$ to $\sim 100 \mu\text{s}$, which allows more fine-grained control of pulse durations and thus integrated pulse amplitude. The highest fidelity single-qubit XY gates reported to date in trapped ion qubits have infidelities of $1.0(3) \times 10^{-6}$ per gate, with a π -pulse duration of $24 \mu\text{s}$ [18]. Because of their longer duration, the spectral content of these control pulses is fairly narrow. Off-resonant coupling to other states is generally negligible, given typical separations to the nearest neighboring hyperfine transitions of $\gtrsim 5 \text{ MHz}$, and sometimes $\sim 100 \text{ MHz}$, depending on the ion species and the magnetic field. Zeeman qubits have no additional neighboring levels and so off-resonant excitation of other levels is not a concern.

For simplicity, pulses are generally rectangular, without shaped rise and fall times, because even the broader spectral content from the sharp pulse edges is far detuned from other transitions. Pulses with durations much longer than $\sim 100 \mu\text{s}$ can become problematic for field-sensitive qubits, if the fluctuations of the qubit frequency due to environmental magnetic field noise start to become comparable to the Rabi frequency $\Omega/2\pi$.

Unlike with superconducting qubits, where each qubit has a dedicated drive line, microwave control fields are relatively uniform over an array of trapped ions, providing only global control. This arises because the ion-ion spacing is much smaller than either the distance to the driving antenna/horn or the free-space wavelength of the control fields. In surface-electrode traps, “regional” control can be achieved through “beamforming” from multiple antennas spaced a suitable distance apart [56]. Individually addressed XY gates for closely-spaced ions can be realized in one of several ways. For magnetic-field-sensitive qubits, an applied magnetic field gradient along the ion string causes each ion to have a separate qubit frequency, so that frequency domain addressing is possible [57]. Alternatively, focused laser beams, or microwave field gradients from near-field electrodes, can be used for

individual addressing by creating differential Rabi frequencies or differential qubit frequencies on multiple ions, among other techniques [58]–[61]. These latter methods can be used on both field-sensitive and field-insensitive (clock) qubits. The full literature for trapped ion qubit individual addressing techniques, including individual addressing using laser beams, is extensive and is not referenced here.

Implementing XY control of single semiconductor spin qubits requires microwave magnetic fields B_1 , applied orthogonal to the direction of B_0 , and typically generated from dedicated on-chip antenna structures near each qubit. An alternate approach uses a global cw microwave magnetic field on all qubits in the array [22]. Individual qubits are then tuned in and out of resonance with the global field via local gate electrode voltages, effectively pulling or pushing the electron wavefunction towards an interface to modify the g -factor [22], or away from the nuclear spin of a donor atom to vary the hyperfine coupling [22], [30].

For S-T qubits, static magnetic field gradients between the two dots that make up the qubit can drive XY rotations. For control of these rotations, nanosecond rectangular pulses are used to separate the two electrons for a time such that they experience different magnetic fields. These gradients are either produced naturally by hyperfine magnetic fields from neighboring nuclear spins, or engineered using micro-magnets on the surface of the semiconductor (the latter is better controlled and gives reduced gate error). Resonant microwave driving at the frequency corresponding to the exchange energy ($\lesssim 1 \text{ GHz}$) can also be used for manipulation of S-T qubits, with the advantage that the drive frequency is much lower than for single-spin qubits [62].

The need for magnetic fields can be alleviated altogether using the three-electron E-O spin qubit. Here, qubit XY control is implemented using modulation of the exchange energy between two of the electrons, proportional to the overlap of their wavefunctions and controllable using time-dependent voltages applied to surface gates, similar to S-T qubits. Exchange between the right-most spin pair drives qubit rotations about the Z axis, whereas exchange between the left-most pair drives rotations about an axis tilted by 120 degrees from the Z axis. Concatenating up to four pulses produces single qubit rotations around any axis. Again, these pulses are typically baseband square pulses with nanosecond rise and fall times. The use of the exchange interaction for qubit XY control reduces microwave complexity, but does so at the cost of increased sensitivity to voltage noise on the control line and charge noise in the material system.

2) Z GATES

In addition to XY gates, a universal gate set requires rotations about the Z axis, which are referred to as Z or phase gates. Referring to Fig. 9, Z gates only affect the qubit phase ϕ . One can perform these gates either virtually, by applying a phase jump to the RF carrier used for subsequent XY gates [63], or physically, either by applying a sequence of two XY gates

whose combined phase values yield a prescribed Z rotation or by detuning the qubit frequency by $\delta\omega_{01}$ for a controlled duration τ , so as to accumulate phase (similar to how phase is adjusted in a phased locked loop during acquisition). Virtual Z gates are attractive since they can be instantaneously applied to baseband IQ envelopes and their accuracy just requires a stable system clock, which is necessary to generate high fidelity XY gates in the first place. However, for systems with more than one qubit, control of the relative phases of qubits is often necessary, which requires physical Z gates.

Minimizing error rates requires carrying out physical Z gates quickly, so it is often preferable that these gates be mediated via a frequency tuning mechanism rather than via multiple XY gates. For superconducting transmon qubits, control of ω_{01} can be enabled by replacing the single JJ of Fig. 6(a) with a two-JJ loop, known as a superconducting quantum interference device (SQUID) (Fig. 7(b)). In this context, the SQUID behaves as a magnetic-flux-tunable nonlinear inductor, so the qubit frequency can be externally controlled via a current bias, analogous to how a voltage controlled oscillator is tuned via a control voltage. Coupling is kept very weak to minimize frequency fluctuations due to noise on the flux-bias tuning line, with typical mutual inductances in the range of just 2 pH. This weak coupling leads to control currents of a hundred microamps or more, and nanoamp-level resolution is required to minimize errors. Digital-to-analog converters (DACs) with at least 14 bits of resolution are typically used to generate these currents.

For gate-defined semiconductor spin qubits, the specific method of producing Z rotations depends on the flavor of qubit, but generally amounts to applying a rectangular pulse to modulate the energy detuning or the exchange interaction between electrons. For single spins, square-shaped pulses can be applied to vary the position of the electron wavefunction, modulating the Zeeman energy for the duration of the pulse. For S-T qubits, turning on the exchange coupling swaps the two spins in the presence of the magnetic gradient, effectively performing a rotation about the Z axis. Similarly, for the three electron E-O qubit, rotations about X, Y , or Z are produced via application of concatenated rectangular pulses to modulate the exchange interaction between left and middle electrons, or right and middle electrons.

Physical Z rotations in trapped ion Zeeman qubits and magnetic-field-sensitive hyperfine qubits (where $\frac{\partial\omega_{01}}{\partial B_0} \neq 0$) could be implemented by changing the magnetic field B_0 to shift the qubit frequency by $\delta\omega_{01}$ for a specified duration τ , yielding a Z rotation by $\delta\omega_{01}\tau$. However, the static magnetic fields in trapped ion experiments are usually generated by highly stabilized current sources or permanent magnets, which are not amenable to rapid field shifting and would provide only global Z rotations. Therefore, physical Z rotations are typically implemented by compiling them to XY gates, applying ac Stark shifts to the ions using detuned laser beams, or applying ac Zeeman shifts to the ions using detuned microwave or rf magnetic fields. Individually addressed Z rotations can be applied using focused laser beams to shift




FIGURE 10. (a) Coupling of linear resonators via a switched capacitance. (b) Coupling of transmon qubits via a tunable coupler. Two qubit gates can be mediated using the qubit and coupler bias currents (i_{Z1} , i_{Z2} , and i_{CPL}).

only specific ions, or by using a gradient of an rf/microwave magnetic field [60], [64]. These techniques have the added benefit that they can be applied to hyperfine clock qubits as well.

B. MULTI-QUBIT GATES

While a library of single qubit gates allowing arbitrary control on the Bloch sphere is required for universal quantum computing, this alone is not sufficient. As described earlier, we also need to create entangled states if we are going to exploit the exponentially large computational space. To do this, we must carry out operations involving multiple qubits, which requires entangling the qubits in a deterministic manner. To understand how one might interact qubits, we can begin by considering the circuit of Fig. 10(a), in which two linear resonators are coupled through a switched capacitor network. Assuming an ideal switch in the open position, the resonators are isolated and the natural frequencies of the circuit are just those of the individual resonators: $\omega_1 = 1/\sqrt{L_1C_1}$ and $\omega_2 = 1/\sqrt{L_2C_2}$. However, when the switch is closed, the resonators pull each other and the natural frequencies of the circuit shift. If each resonator was initially oscillating at its respective resonant frequency, the oscillation dynamics would shift from a distinct frequency at each node to the joint modes of the coupled resonator system (which will be shifted from the bare frequencies. In the case where the coupling is weak, we can think of the coupling as a perturbation applied to each of the isolated resonance frequencies. As such, if we energize one of the resonators and pulse the switch closed for a short period of time, the oscillation in the resonator will acquire a phase shift due to the temporary shift in resonance frequency.

Now, imagine both of these resonators have a nonlinear relationship between their isolated natural frequencies and the amplitude of oscillation (similar to a qubit). In this case, if we perform the same experiment where we energize one of the resonators, pulse the coupling on, and look at the acquired phase, the result will depend upon the state of the second

resonator, since the degree to which it pulls the first resonator is state dependent. This behavior describes one way the basic interactions required to perform entanglement generating two-qubit gates can be carried out.

Two-qubit gates require a mechanism for interacting qubits in a deterministic manner, and many different approaches have been proposed and demonstrated. For instance, interactions between superconducting qubits can be engineered at the circuit level by introducing either static [65] or tunable couplings [7], which may be implemented inductively or capacitively. One approach is shown in Fig. 10(b), where a tunable coupler allows interactions between a pair of frequency tunable transmons [66]. The coupler itself consists of a transmon—which serves as a frequency tunable LC resonator—that is capacitively coupled to each of the qubits. With this structure, it is possible to realize coupling strengths ranging from completely off to tens of MHz. One can engineer a wide range of two-qubit gates through proper design of the three current bias waveforms for this circuit [67]. For instance, if the qubits are tuned into resonance and the coupling is enabled, a single excitation will oscillate back and forth between the qubits (that is, oscillation will occur between the probability amplitudes of the $|01\rangle$ and $|10\rangle$ states), and by properly setting the gate duration it is possible to engineer a gate where α_{01} and α_{10} are swapped.

While two-qubit gates for superconducting qubits can be carried out through control of the qubit frequencies as described above, they can also be performed without changing the qubit frequency, using an additional microwave drive tone instead. Such approaches are necessary in architectures that use fixed-frequency qubits (see for example Ref. [42]). One example of an all-microwave two qubit gate is the cross-resonance (CR) gate [68], in which one of a pair of reactively-coupled qubits is driven by a microwave tone at the qubit frequency of the other qubit. With the appropriate drive amplitude and duration along with the addition of a single qubit gate applied to each of the qubits, the CR gate can be used to implement a CNOT operation, which swaps α_{10} and α_{11} while leaving α_{00} and α_{01} unchanged. The gates described here are just two of a wide assortment of gates that can be applied to superconducting qubits; for further discussion of two-qubit gates in superconducting qubits, we refer the reader to Ref. [39].

Two-qubit entangling gates in semiconductor spin qubits are generally implemented via the Heisenberg exchange interaction between electrons on neighbouring dots of an array [29]. The charge dipole associated with the two-electron system opens the prospect of coupling S-T qubits capacitively, since the relative spin orientation of one qubit can lead to charge rearrangement that effectively gates another qubit [69]. The advantage of exchange coupling is its controllability, modulating the tunnel coupling between two adjacent quantum dots for a controlled amount of time. The evolution of the two-qubit system then depends on the wave-function overlap of the electron states, resulting in the physical exchange of the electron positions to execute a $\sqrt{\text{SWAP}}$ entangling gate. A

major technical challenge for exchange coupled qubits is the sensitivity of the tunnel rate to gate voltage and the requirement that electrons must be brought within nanometers of each other for coupling. This latter aspect leads to crowding of gate electrodes and challenges for crosstalk mitigation. Alternative coupling schemes making use of intermediate electron states [32] are presently an active area of research. A further fruitful direction is to couple remote S-T qubits via a cavity resonator, following similar approaches to superconducting qubits [70]. To date, two-qubit gates using E-O qubits have not been demonstrated, although qubit coupling schemes are likely to be similar to S-T qubits.

Trapped ion hyperfine and Zeeman qubits have negligible direct interaction with each other due to very weak spin-spin coupling. However, the motion of multiple ions in a single trapping potential is very strongly coupled. As a result, almost all entangling gates carried out between trapped ion qubits are realized using the quantum motion of the trapped ions as an intermediary “bus”. An effective ion-ion interaction is realized by coupling the trapped ion spin to its motion in an appropriate way, using external control fields. Typically this is done with laser beams, but it is also possible to do using rf and/or microwave fields. Crucially, spin-motion coupling requires a spatial gradient of the control field over the spatial extent of the ions’ quantum mechanical zero-point motion in the trap, typical ~ 10 nm. The magnetic field gradient of free-space microwaves near ω_{01} (usually a few GHz) is very small over this length scale. However, microwave magnetic fields with negligible gradient strength can be combined with additional static [71], [72] or few-MHz [64], [73] magnetic field gradients to produce the desired spin-motion coupling. Magnetic field gradients at microwave frequencies near ω_{01} , or near resonance with the ion motional frequency, can also be used for spin-motion coupling and entangling gates [10], [74]. These gate protocols are often carried out in surface-electrode traps, where larger near-field magnetic field gradients can be generated [46], [75], [76].

A number of experimental demonstrations of high-fidelity microwave-based entangling gates between ions have been carried out [61], [77]–[80]. In general, these gates are slower (\sim ms duration) than laser-based entangling gates, which can be performed in tens or hundreds of μs (and some as fast as a few μs [81]). However, the fidelities reported for microwave-based gates, with errors in the few 10^{-3} range per gate, are competitive with the fidelities of laser-based gates. In addition, laser-based gates have fundamental fidelity limitations due to off-resonant scattering from excited electronic states in the ion [82], making microwave-based gates, which do not have this fundamental limit, an appealing alternative.

C. HARDWARE FOR QUANTUM STATE CONTROL

High-level control requirements for each of the three qubit technologies are summarized in Fig. 11. While the details vary greatly among the technologies, some form of pulsed RF waveforms and/or baseband control signals are required to run a quantum processor. An exemplary control system for a pair

	Microwave Control	Baseband Control
Superconducting (Transmon)	1Q XY gates, 2Q gates Carrier: 4-8 GHz Pulse duration: 10-30 ns π -pulse P_{av} : -80 to -60 dBm Shaped envelope	1Q Z gates, 2Q gates 0,01-1mA static/pulsed Pulse Duration: 10-500 ns Resolution: ~nA
Semiconductor Spin	Single spin Q: 1Q XY gates Carrier: 0.1-50 GHz Pulse duration: 10 ns to 1 μ s π -pulse P_{av} : -60 to +0 dBm Shaped envelope	Single spin Q: 2Q gates S-T Q: XY gates, 2Q gates E-O Q: XY gates, 2Q gates μ V-mV level signals Pulse Duration: ns-ms 1 ns rise/fall
Trapped Ion	1Q XY & Z gates, 2Q gates Carrier: 5-20 MHz, 1-12.6 GHz Pulse duration: 1-500 μ s π -pulse P_{av} : 0 to 45 dBm Rectangular envelope	Qubit state control typically not performed at baseband

FIGURE 11. Summary of microwave/baseband control requirements for each of the qubit technologies. Abbreviations: Q–qubit, P_{av} –available power at qubit drive port.




FIGURE 12. Simplified control hardware for a two-qubit transmon-based quantum processor with tunable coupler.

of superconducting qubits connected via a tunable coupler appears in Fig. 12. Each XY control signal is generated using single-sideband mixing, with the complex envelope generated using a pair of high-speed DACs. Alternatively, the RF XY signals are sometimes directly generated using high speed DACs [83], [84], obviating the need for analog mixing. The XY signals are heavily attenuated to suppress thermal noise at the qubit drive port. An additional three DACs generate the broadband control signals required to drive the Z control lines and coupler bias port. These DACs feature 14 bits of resolution, as required for the frequency control of the qubits and coupler. The digital waveforms are generated using a field-programmable gate array (FPGA), which is configured to orchestrate quantum algorithms. Architectures such as this are extensible to hundreds of qubits and widely in use among

the quantum computing community. Part of the reason that the field of quantum computing has experienced rapid growth over the past ten years is that components required to build control systems such as this have become commodity items, thanks to the wireless communications revolution.

V. MEASURING THE STATE OF A QUBIT

To behave quantum mechanically, as required for quantum computation, qubits must be well-isolated from sources of noise or dissipation in their environment. We have described this requirement thus far by stating that Q_d , Q_i , and Q_c must be large ($> 10^6$, and ideally even higher). However, at some point in a quantum algorithm, the state of the qubit must be measured, a task which necessarily requires coupling the qubit to the outside world strongly enough that the qubit can influence the state of the measurement apparatus and thus allow us to determine the qubit state. This coupling is characterized by a quality factor Q_m which describes the effective dissipation bath seen by the qubit due to coupling with the measurement apparatus.

Because the energy difference between the two qubit states $\hbar\omega_{01}$ is very small for microwave-frequency qubits, some form of amplification must be used to convert it into a classical-level signal strong enough to be digitized and analyzed to read out the state of the qubit. Multiple stages of amplification are used; as with standard low-noise microwave receivers, the overall performance is most sensitive to the first amplification stages.

Superconducting qubits, and increasingly semiconducting qubits as well, typically use a technique called dispersive readout to measure the qubit state [85]–[89]. This approach relies on coupling the qubit to a microwave-frequency superconducting resonator. The resonator frequency ω_r is detuned from ω_{01} by an amount $\Delta \gg g$, where g is the qubit-resonator coupling strength. In this so-called “dispersive regime” of the coupling, the qubit and resonator modes do not exchange energy, but the frequency of the resonator is shifted by an amount 2χ , called the “dispersive shift”, depending on the qubit state. The resonator is overcoupled to a feedline, giving a resonator linewidth κ which is ideally $\sim 2\chi$ (typically $\kappa \sim 2\chi \sim 2\pi \times 5$ MHz). The resonator is typically probed in reflection, as shown in Fig. 7(c), with a probe tone near ω_r . The reflected signal will depend strongly on whether the resonance frequency has been shifted, which in turn indicates the qubit state. Thus the measurement of the qubit state can be realized by measuring S_{11} of the readout resonator at a single frequency near ω_r . For strongly overcoupled resonators, $|S_{11}| \approx 1$, but $\arg(S_{11})$ will be qubit-state-dependent near ω_r . One can think of this as a binary phase-shift-keyed signal, whose symbols (typically separated by $\approx 180^\circ$, but potentially separated by a smaller angle) correspond to the two qubit states. We note that this process also constitutes a kind of amplification: the single-photon qubit energy $\hbar\omega_{01}$ has been turned into a phase shift on many cw probe photons.

The challenge here lies in the fact that the probe tone must be very weak, so as not to scramble the qubit state after

initially collapsing it to $|0\rangle$ or $|1\rangle$ [90], [91]. Otherwise, the resonator frequency would jump back and forth along with the qubit state, and the readout signal would stop providing information about the state to which the qubit initially collapsed. The limit on the probe tone amplitude depends on the vector difference in the IQ plane between the reflected probe tone corresponding to state $|0\rangle$ and the reflected probe tone corresponding to state $|1\rangle$ —in other words, the size of the phase-shift-keyed signal. For superconducting qubits this difference signal should generally be kept below ~ 300 nV_{peak}, which corresponds to a signal power of around -120 dBm. For spins, probe powers up to -80 dBm have been used.

This weak probe tone must be amplified, but the amplifier noise must be small enough not to drown out the signal. While one can average in time (in other words, decrease the resolution bandwidth) to remove amplifier noise and recover the weak signal, we must perform our measurement in a time $\ll T_1$, so that the qubit state does not decay during readout and thus corrupt the readout result. Practically speaking, this means that readout must be performed in several hundred nanoseconds, or equivalently that the resolution bandwidth must be $\gtrsim 10$ MHz. The duration of measurement also poses a limit to the clock rate of a quantum algorithm or error correcting code.

The solution is to use an ultra-low-noise amplifier as the first stage of the receiver, such that even the weak readout signal can be amplified with signal-to-noise ratio well above 1. Superconducting parametric amplifiers based on Josephson junctions as nonlinear elements [92] can reach the quantum limit for noise, where the amplifier noise temperature T_N at frequency ω is equal to $\hbar\omega/2$. For 6 GHz signals, this corresponds to $T_N = 144$ mK. Dispersive qubit readout using superconducting parametric amplifiers was first implemented a decade ago [93], and readout fidelities of $\sim 99\%$ have been reported [94]. Progress in superconducting parametric amplifiers in the past 15 years has been driven by quantum computing applications; we discuss this in more detail in Section VII.A.

As discussed above, it is important that Q_m be high, so that the measurement process not couple the qubit to a lossy environment that would cause T_1 decay. The readout resonator acts as a bandpass filter between the qubit and the lossy (real) impedance of the readout transmission line; since the resonator is far detuned from the qubit, this suppresses coupling between the qubit and this source of loss. However, for some qubits this coupling is still the dominant source of loss. A standard way to mitigate this loss is to place a bandpass filter (usually a second resonator) between the readout resonator and the transmission line [94], [95]. This type of filter is known as a *Purcell filter* and can be used to boost Q_m by an additional two orders of magnitude or more, depending on the design parameters. The bandwidth of the Purcell filter can be much larger than κ , so that it does not affect the speed of the readout.

It is also possible to read out a superconducting qubit by turning on a strong coupling to a microwave photon counter.

If the qubit were in state $|1\rangle$, it would decay to $|0\rangle$ and emit a photon of energy $\hbar\omega_{01}$ into the photon counter. It is challenging in practice to realize a microwave photon counter because the single photon energy is so small in the few-GHz regime, but such a device (which is essentially a modified superconducting phase qubit) has been demonstrated with readout fidelity as high as 98.4% [96], [97].

For trapped ion qubits, the very weak coupling of the qubit state to microwave fields (much higher Q_d than the other technologies) means that it would be very difficult and inefficient to extract information from the ion at microwave frequencies. However, trapped ions possess optical transitions, allowing easy extraction of optical photons, which have the added benefit of being easily detected with low background noise by room-temperature single-photon counters.⁶

Trapped ions are generally read out using the so-called electron shelving technique [98], where the probability amplitudes in states $|0\rangle$ and $|1\rangle$ are mapped to two suitable states $|b\rangle$ and $|d\rangle$ using coherent control pulses of the same types used for qubit manipulation (microwave pulses and/or laser pulses). When illuminated with a laser beam of appropriate wavelength and polarization, an ion in state $|b\rangle$ will fluoresce, absorbing photons and re-emitting them in all directions. The laser beam is chosen to drive a so-called cycling transition, where the ion is excited from the state $|b\rangle$ and then emits a photon, always returning to $|b\rangle$ after the emission. This enables repeated rounds of excitation and emission. It is sometimes necessary to use multiple laser beams to “close” this cycle. In contrast, an ion “shelved” in the state $|d\rangle$ (chosen such that all transitions out of $|d\rangle$ are far off resonance with the readout laser) will not interact with the laser beam and thus will not give off fluorescence photons. By collecting a fraction of the fluorescence photons with an imaging objective and counting them with a single photon counter, it is possible to distinguish between a fluorescing “bright” ion in state $|b\rangle$ and a “dark” ion in state $|d\rangle$, as long as the mean number of photons counted for bright and dark ions is sufficiently different. In practice, a readout duration of several hundred microseconds typically gives tens of counts for a bright ion, and $\sim 1 - 4$ counts for a dark ion, although the duration and count rates can vary by an order of magnitude or more depending on the specifics of the setup. Readout fidelities as high as 99.99% have been demonstrated using this technique [99].

Electron shelving readout implements two forms of amplification. First, mapping into the states $|b\rangle$ and $|d\rangle$ and scattering a single photon can be thought of as turning the qubit energy difference $\hbar\omega_{01}$ into the energy of an emitted UV or visible photon, which is typically between 10^5 and 10^8 times larger for hyperfine or Zeeman qubits. Secondly, the cycling transition allows up to $\sim 10^6$ such photons to be scattered during the readout operation, giving further gain.

⁶This is primarily because of their high energy relative to the available thermal energy ($\hbar\omega \gg k_B T$ for optical photons at $T = 300$ K).

VI. OTHER APPLICATIONS OF MICROWAVE TECHNOLOGY

Beyond the direct use of microwave signals and techniques for qubit state control and measurement, there are many other critical applications of microwave and rf technology in quantum computing. This section details several of the essential supporting roles played by microwave technology, beyond direct control and measurement of qubits.

A. RF FOR ION TRAPPING

Trapped ion qubits for quantum computing applications are almost always confined using a combination of radio-frequency and static electric fields in a so-called Paul trap [10] (Penning traps, which use a combination of static magnetic and electric fields to confine ions, can also be used for quantum computing but are more typically employed for quantum simulation or precision spectroscopy experiments [100]–[102]). The rf electric fields for trapping are generated by applying rf voltages to specific electrodes of the trap, such that an rf electric field quadrupole is formed in vacuum at a distance of typically 30 μm to several mm from the electrodes. The oscillating quadrupolar rf fields provide an effective confining potential for ions. We refer the reader to Refs. [10], [11] for further details on Paul traps.

The frequency and amplitude of the applied rf necessary for stable trapping of ions depend on the physical dimensions of the trap, the species of ion, and the desired strength of the confinement. Larger traps, heavier ions, and stronger confinement all require larger rf voltages on the trap electrodes. Reducing the rf frequency enables trapping with lower rf voltages, but if the frequency is reduced too much then the trap will become unstable and the ions cannot be held reliably. Typical rf Paul traps for quantum computing applications operate at rf frequencies between 20 MHz and ~ 100 MHz, and with peak rf voltages on the trap electrodes between 10 V and ~ 500 V.

These combinations of frequency and amplitude are generally quite difficult to achieve with direct driving, so the rf electrodes are typically incorporated into an rf resonant structure, which provides rf voltage step-up and filters noise from the rf drive electronics that could excite the ion motion. Many different classes of resonators have been used, including helical resonators [103], [104], coaxial resonators [105], and lumped element resonators [106], [107]. These resonators typically have loaded quality factors of several hundred, although some cryogenic resonators can approach 10^4 . Toroidal transformers on ferrite cores, operating as flux-coupled impedance-transforming ununs, can also be used to provide voltage step-up [108].

Qubit-qubit coupling for trapped ions relies on the use of the ions' shared motional degree of freedom as a coupling "bus". Multi-qubit gates typically have durations τ_g from ~ 10 μs up to several ms [12], with high-fidelity microwave-based gates typically in the $\sim\text{ms}$ range. To achieve high gate fidelity, the motional frequency must be stable to $\ll \frac{1}{\tau_g}$. In practice, this means that the ion motional frequencies, typically a few MHz, should be stable at the ~ 10 Hz level or below, or better

than a part in 10^5 . In a typical linear rf Paul trap, 1/3 of the motional modes (called "axial" modes) have frequencies determined purely by static potentials applied to trap electrodes and 2/3 of the modes (called "radial" modes) have frequencies that depend linearly on the rf voltage applied to the trap electrodes. Many multi-qubit gate schemes use radial modes, so achieving the required motional frequency stability of ~ 10 ppm or better means that the rf voltage on the trap electrodes must be stable at the ~ 10 ppm level as well. This requires high stability of the resonator frequency and quality factor, high gain stability for the rf amplifier driving the resonator, and low amplitude and phase noise for the rf generator. If the resonator Q is too high, this can cause increased voltage fluctuations at the trap through FM-to-AM conversion, giving some intuition for why trap rf resonators are not designed simply for maximum Q . Recent work in the field has demonstrated rf limiting amplifiers [109] and active feedback methods for rf amplitude stabilization [109], [110] which can meet the desired ~ 10 ppm amplitude stability of the rf voltage at the trap.

B. APPLICATIONS TO LASER SYSTEMS FOR QC

Numerous qubit technologies, including trapped ion qubits, neutral atom qubits [111], [112], and optically-active defect centers in solids [113], [114], rely on laser-based control and readout methods. These laser systems use microwave technology for a number of tasks. Most laser beams for quantum information experiments need to be pulsed on and off with microsecond rise/fall times, too fast for a mechanical shutter. In addition, sometimes a laser beam with amplitude modulation, rapidly tunable frequency, or multiple frequency components in a single beam is desired. For these tasks, acousto-optic modulators (AOMs) [115] and/or electro-optic modulators (EOMs) [116] are used. An EOM can be thought of like a microwave mixer, except the LO port uses an optical signal while the IF port accepts an rf or microwave signal; signals at the IF port modulate the optical LO and give rise to sidebands on the optical output. AOMs are more akin to single-sideband mixers. They perform frequency-shifting of the laser light passing through them, but can also be used for amplitude modulation (including on/off switching) of laser beams. The "carrier leakage" and "spurious sideband" laser beams have different directions of propagation from the desired "sideband" at the AOM output, enabling them to be filtered out spatially. Both AOMs and EOMs require rf or microwave signals at around +30 dBm to function; generally AOMs are driven at frequencies from ~ 50 to ~ 500 MHz, with 10s of MHz bandwidth, while EOMs can accept drive tones from near dc up to many GHz, and have widely varying bandwidths depending on the choice of resonant rf/microwave enhancement circuits inside.

In addition to modulation of laser light, microwave technology is also important for stabilization of the lasers themselves. A famous method for locking a microwave oscillator to a stable reference cavity using modulation sidebands, due to

Pound [117], was extended to operation with laser oscillators and stable optical references (cavities or spectral lines) by Drever, Hall, and coworkers [118], and is known in the laser community as Pound-Drever-Hall (PDH) locking. PDH lock circuits, which operate at rf/microwave frequencies, are ubiquitous in laser systems. There are also a wide variety of servo loops for laser amplitude stabilization which rely on modulating the amplitude of the rf drive to an AOM as the feedback signal. Mode-locked lasers, which emit periodic short (\lesssim ps) pulses of light with an extremely stable repetition rate, forming a frequency comb, are used in trapped ion applications to drive stimulated Raman transitions between states with very large detuning, such as the 12.6 GHz hyperfine qubit states in $^{171}\text{Yb}^+$ [119]. For optimum performance, the repetition rate of the mode locked laser (typically ~ 100 MHz, but potentially as high as a few GHz) must be stabilized, which requires implementing a phase-locked loop between a stable reference oscillator at the desired repetition rate and the signal from the laser pulse train on a fast photodiode.

VII. MICROWAVE INNOVATIONS FROM QUANTUM COMPUTING

Quantum computing relies heavily on microwave technologies already developed for other applications. However, quantum computing also requires operation in new performance regimes, and has inspired the development of novel microwave technologies and systems to meet those challenges. We describe two major advances in microwave technology that have arisen from quantum computing research: quantum-limited microwave amplifiers, and cryogenic non-reciprocal microwave devices, including chip-scale non-reciprocal devices.

A. QUANTUM-LIMITED AMPLIFIERS

As described in Section V, the readout of superconducting qubits and semiconducting qubits involves microwave signals so weak that they are not far above the quantum noise floor of half a photon per unit bandwidth, $S(\omega) = \hbar\omega/2$. Faithful amplification of these signals requires amplifiers with noise performance at or near the quantum limit, which at 6 GHz corresponds to a noise temperature of 144 mK.

To achieve this noise performance, the field has turned to superconducting parametric amplifiers. Parametric amplifiers rely on a nonlinear element or elements whose parameters (inductance or capacitance) are modulated in time by a strong pump. This modulation transfers power from the pump into other modes (known as the signal and the idler), coherently amplifying the energy in those modes [92], [120].

Before the advent of high-electron-mobility transistor (HEMT) amplifiers, microwave amplification was sometimes accomplished using parametric amplifiers based on varactor diodes or inductors with saturable cores [121]. Parametric amplifiers based on the Josephson effect in superconductors have been studied intermittently since the 1960s, but the field was revitalized by work in the late 2000s on Josephson parametric amplifiers with near-quantum-limited noise performance

and bandwidths of ~ 10 MHz [122]–[125]. The viability of these amplifiers for high-fidelity qubit readout was demonstrated shortly thereafter [93], and they were rapidly adopted as the state of the art [94], [126]. Subsequent work has used impedance engineering to increase bandwidths to the \sim GHz range [127]–[129], and more complex designs with many Josephson junctions have shown substantial improvement in saturation powers [129], [130]. Most designs operate in reflection, typically requiring bulky circulators to separate the output signal from the input signal, and to protect the qubit circuit from the strong pump tone, so transmission-mode parametric amplifiers providing directional, non-reciprocal amplification have been developed [130]–[133]. A considerably more extensive review of the literature than is provided here can be found in Ref. [92].

The bandwidth, saturation power, and noise performance of current state-of-the-art parametric amplifiers allows simultaneous readout of multiple qubits by frequency multiplexing their readout resonators, which share a common feedline to the parametric amplifier. The gain and noise performance of these parametric amplifiers are such that they set the overall receiver noise for the amplification chain (the rest of which consists of ultra-low-noise cryogenic and room-temperature microwave transistor amplifiers) out to the room temperature demodulation and digitization circuitry.

B. NON-RECIPROCAL DEVICES

Microwave circuit elements exhibiting non-reciprocity are currently used heavily in quantum computing, mostly in the amplification of readout signals from qubits. In this context they serve two main purposes. First, they function as isolation devices that prevent noise originating in the readout amplification chain from impinging on the qubits. Their second use is in the context of parametric amplifiers (see above) that operate in reflective mode. Here, circulators are configured to separate input and output signals as well as to isolate the qubits from the pump tones that supply energy to the amplifier. To achieve non-reciprocity, traditional microwave circulators exploit ferromagnetic materials. These devices are necessarily large and bulky components, since they make use of interference effects that occur over a length scale comparable to the microwave wavelength. Further limitations include their typically narrow-band performance, insertion loss, limited non-reciprocity, and variation in specifications or failure when operated at cryogenic temperatures or in large magnetic fields. From the perspective of scaling up to the number of readout channels needed for a large quantum system, the footprint alone of conventional circulators is a significant barrier to the development of tightly integrated systems.

Recently, efforts have focused on realising miniaturized devices that exhibit nonreciprocity. One avenue uses active devices, specifically Josephson parametric amplifiers designed to achieve directional gain and reverse isolation while maintaining quantum-limited noise performance [130]–[132],




FIGURE 13. On-chip microwave circulator based on the quantum Hall effect. Microwave excitations travel at a velocity 1000 times slower than in vacuum, enabling an equivalent reduction in the signal wavelength and device footprint [134].

[135], [136]. Ref. [133] provides an in-depth review of nonreciprocal active circuits with quantum-limited noise. Such devices have recently been demonstrated for direct, high-fidelity readout of superconducting qubits without external circulators [137]–[139].

Miniaturized passive circulators also hold promise for quantum computing applications. For example, Fig. 13 shows a 3-port circulator device based on the inherent nonreciprocity of the quantum Hall effect [134]. Here, the size of the device is 1/1000th of the free-space wavelength of the microwave signals it handles. This small size is achieved by exploiting the chiral, “slow-light” response of a two-dimensional electron gas in the quantum Hall regime. For an integrated GaAs device with 330 μm diameter and 1 GHz center frequency, a nonreciprocity of 25 dB is observed over a 50 MHz bandwidth. Furthermore, the nonreciprocity can be dynamically tuned by varying the voltage at the port, an aspect that may enable reconfigurable passive routing of microwave signals on chip. The forward transmission of this particular device was limited to -20 dB due to the impedance mismatch between 50 Ω environment and the device impedance, which is set by the resistance quantum $R_K \approx 25$ k Ω . However, improved impedance matching design could be used to increase the forward transmission [140]. Although the quantum Hall device requires a large external magnetic field, magnetically-doped topological insulator materials offer a further route to realizing non-reciprocal devices without the need for such external magnetic fields [141]. Beyond non-reciprocal behavior, the underlying physics of these systems suggests their use as compact and tunable delay lines, microwave interferometers, and low-loss, high impedance transmission lines [142].

For future quantum computing systems, miniaturized, on-chip non-reciprocal microwave devices of the kinds described above are likely to play a key role.

VIII. MICROWAVE CHALLENGES IN REACHING THE FULL POTENTIAL OF QUANTUM COMPUTING

In order to realize the full promise of quantum computing, it will be necessary to build systems capable of executing QEC protocols. In general, these protocols use ensembles

of error-prone physical qubits to implement “logical qubits,” which have reduced rates of error in comparison to their constituent physical qubits. The degree to which errors are suppressed depends both upon the error rates of the physical qubits and the degree of redundancy. However, for such a protocol to work in the first place, the physical qubit error rates must be on average below a threshold, which depends on the specifics of the QEC protocol. There are many different QEC schemes [48], [49], with physical error thresholds ranging from $\sim 10^{-4}$ up to about 1 %. Once the physical error rates are below the threshold, the degree of redundancy (the number of physical qubits) required to achieve a specific logical error rate scales logarithmically with further reduction in the physical error rates. A practical target of exceeding the threshold by an order of magnitude (0.1% physical qubit error rates for some codes) appears feasible, but still requires about 1,000 physical qubits per logical qubit. Thus, it is estimated that realizing a practical system with 1,000 logical qubits (far below the number required to break RSA encryption) will require building a million-qubit-scale system [51]. There are numerous microwave-related challenges associated with building such a computer. Here we provide a discussion of some of these challenges.

Today’s quantum processors are controlled and measured using systems that mostly comprise room temperature electronics, with high-performance microwave interconnects used to interface to the quantum processor. Let us begin by considering the feasibility of scaling this approach to the million qubit level. It will be helpful to start by quantifying the number of control and measurement channels required for a brute force scaling approach. For a large-scale quantum processor implemented using superconducting transmon qubits arranged in the architecture of [7], processor operation requires one XY channel and three Z control channels per qubit. In addition, one readout channel is required for every ~ 10 qubits. In the case of semiconductor spin qubits a handful of dc bias wires per qubit are needed to define tunnel barriers and chemical potentials. Depending on the flavor of spin qubit at least one high bandwidth channel is required per qubit, with additional channels, shifted in phase, often used to cancel crosstalk. Readout of spin qubits based on rf reflectometry requires a single wideband line per ~ 10 qubits using frequency multiplexing. Trapped ion systems generally rely on their long qubit coherence times, as well as the ability to reconfigure ion and/or laser beam positions in the trap, to perform logic operations on different subsets of qubits in series, rather than performing simultaneous operations on all qubits. This means that the scaling of microwave control hardware does not have a precise relationship to the scaling of qubit numbers. However, to avoid excessive slowing of computations, the serialization would ideally be at most a factor of tens, not hundreds or larger.

Given these constraints, it does not appear feasible to scale existing technology to the level demanded by large-scale error-corrected quantum computers. For instance, if one were to build a quantum control and measurement system for a

million transmon qubits using present Google technology [7], it would occupy over 15,000 m² of floor space and dissipate about 40 MW (not including the dissipation of the amplifiers that would be required to compensate for loss in the cable runs).

A. SCALING OF INTERCONNECTS

Beyond the sheer volume of the quantum control and measurement electronics, one must also consider the feasibility of connecting the electronics to a quantum processor. Here, the fundamental constraint is going to be the ratio of the required number of interconnects to the chip surface area, which sets the required interconnect density. For large-scale quantum processors, this simplifies to the ratio of the number of I/O signals required per qubit to the qubit pitch squared (assuming a 2D array). For superconducting qubits, which are arranged on a pitch of about 1 mm and require about 4.2 lines per qubit, an interconnect density of about 4.2 lines per square millimeter is required; while this is believed to be feasible, it will require significant microwave engineering to deliver these signals while maintaining low crosstalk and avoiding dissipation. For instance, a million qubit quantum processor will require about 3 million Z control lines, and to keep heating due to the $\approx 50 \mu\text{A}$ static Z currents to below 10 μW (necessary because of the limited available cooling power at mK temperatures), the contact resistance between the qubit package and the cable assembly must be kept to the m Ω level. Developing qubit packaging techniques that enable high coherence and low crosstalk while achieving this specification will require significant research. Superconducting interconnects, which offer near-lossless and near-dispersionless electrical propagation [143] while providing very low thermal conductivity, can be used to transport signals from 10 mK to 4 K. However, achieving the required interconnect density between 4 K and room temperature (where superconducting interconnects are not an option) may be challenging due to dispersion and losses, both of which distort control waveforms. As such, the control and measurement system may have to reside partially or entirely at 4 K. Regardless, considerable work is required to develop interconnect systems between 4 K and 10 mK which offer reproducible performance and the stringent cross-talk performance that is required for a large-scale quantum computer.

For semiconductor spin qubits, the nanoscale qubit dimensions make interfacing a large scale quantum processor to an external control system impractical without a large degree of multiplexing [144], [145]. Consider, for instance, that as many as 10 wires are required within the 100 nm \times 100 nm footprint of a spin qubit. This geometric I/O bottleneck motivates the integration of cryogenic electronics with the qubit platform to handle signal generation and multiplexing [146] without needing large cable assemblies that carry signals to room temperature electronics [147]. The power dissipation of these classical electronics can be significant, however, leading to heating of the qubits and a degradation in fidelity if tightly integrated in a monolithic configuration.

Two distinct approaches are currently being pursued to address this trade-off between the I/O bottleneck and qubit temperature. The first involves relocating the spin qubit platform and integrated control circuits to higher temperatures (near 1 K), where substantially more cooling power is available by using pumped helium-4 [27], [28]. However, this approach causes reductions in qubit fidelity, eventually requiring orders of magnitude more noisy physical qubits to encode a logical qubit. The increased qubit count then requires additional I/O and control electronics that may effectively cancel the gains from operating at higher temperature. An alternate approach involves operating the control sub-systems at the same millikelvin temperature as the qubits, but on a separate chip that is thermally decoupled [148]. Taking advantage of the high impedance (open circuit) nature of gate electrodes and leveraging lithographically defined chip-to-chip interconnect strategies, this millikelvin approach alleviates the power dissipation otherwise required to drive low impedance cables between temperature stages.

Large-scale trapped-ion quantum computers will likely operate cryogenically, in the 4 K to 10 K temperature range, because of superior vacuum pressure (enabling ions to remain trapped for much longer times) and reduced electric field noise (improving motional coherence and entangling gate fidelities). The much larger cooling powers available at these temperatures make the interconnect problem less formidable, although still not easy. Because many ions can be addressed with a single microwave control line, a relatively small number of control lines will be needed compared to the other qubit technologies. However, because the control signals have much higher power, on-chip dissipation may present an issue. Ion traps rely on a very large number of static or slowly-varying control voltages for defining trap potentials and moving ions around the trap, with roughly 10 times as many control voltages as potential wells (each of which may hold single or multiple ions) in the trap. However, these lines draw zero or minimal current and do not require high bandwidth (< 1 MHz is generally ample), so they can be made from very thin, low-thermal-conductivity wires. Research into generating these voltages with DACs fabricated in the trap substrate itself is being pursued by some groups [149].

B. SCALING OF CONTROL SYSTEMS

Today's quantum computers are essentially research devices with individual qubit performance just approaching the edge of what is required to implement QEC protocols. Developing control protocols for use on these prototype systems has necessitated the use of flexible, high-speed arbitrary waveform generators so that a researcher can quickly test new control paradigms without developing new hardware. While this is a logical approach for operating today's relatively small-scale quantum processors, the cost, size, and power associated with scaling this approach to run a million-qubit quantum processor motivate the development of a more optimized approach.

Several research groups are currently investigating the integration of quantum control circuits, targeting operation

at a physical temperature of 4 K [150]–[155]. However, the requirements for these systems are stringent and co-optimization of the classical controller and the quantum processor will likely be required. For a device thermalized at 4 K, the power consumption will be limited to well below 1 mW per control channel, and this tight power budget must be met without increasing control errors. For *XY* controllers, strategies must be developed to minimize the number of microwave carriers required, mitigate crosstalk, optimize pulse waveforms, and maintain phase coherence among the ~ 1 million *XY* control channels. The same considerations, albeit with somewhat relaxed power restrictions, are also relevant for room-temperature control electronics.

Similar challenges are present for the development of the scalable baseband waveform generators required for *Z* control. In particular, losses and reflections along the interconnects between the *Z* controllers and the qubit/coupler control ports cause distortion in the same way that non-return-to-zero signals experience distortion when traveling over a backplane. In today's systems, the long cables between the room temperature electronics and the quantum processor lead to settling times that can be on the order of microseconds [147], [156]. To compensate for these long settling times, this response is often characterized and deconvolved from the transmitted *Z* signal, similar to the use of pre-emphasis in wireline applications. While the use of bipolar *Z* pulses reduces the impact of long-term settling, the short time-scale system response must still be compensated [157]. As such, research will be required to architect *Z* control systems that are both precise and low power.

C. SCALING OF READOUT

Scaling of readout is also of critical importance and microwave innovations will certainly be needed. For superconducting and semiconductor spin qubits using dispersive readout, scalable amplification chains achieving near-quantum-limited noise performance are required. Assuming a frequency domain multiplexing factor of $10\times$, a million-qubit quantum controller will require 100,000 readout channels. In today's systems, each readout channel contains a parametric amplifier and several circulators (up to 5) thermalized to 10 mK, a HEMT-based low noise amplifier thermalized to 4 K, and further amplification and digitization at room temperature. In order to scale to the levels required to implement a practical error-corrected quantum computer, each of these technologies will have to be optimized for manufacturability, cost, performance, and size. For instance, the semiconductor low noise amplifiers used at 4 K (e.g., [158]) achieve excellent noise temperatures, but are hand assembled, hand tested, and dissipate about an order of magnitude too much power for use in a large-scale system. As with the case of control systems, it is necessary that scalable readout systems be integrated to the maximum extent possible. Amplifiers that can be mass manufactured in a silicon technology platform and achieve similar performance to today's HEMT LNAs while dissipating $\ll 1$ mW should be developed so that electronics

currently at room temperature can be directly integrated with the low noise amplifier and thermalized to 4 K. While SiGe processes appear to be a promising technology for meeting these goals [159], [160], research is still required to determine if it is possible to achieve both the performance and repeatability required using SiGe technologies. Additionally, work is required to minimize the volume of the microwave electronics that must be thermalized to the 10 mK stage (i.e. the parametric amplifiers and circulators).

D. OTHER MICROWAVE CHALLENGES

Beyond the classical to quantum interface, there are significant microwave challenges in the design of the quantum processor itself. For instance, considering superconducting qubits, as the dimensions of the quantum processor grow, designers will have to rely more heavily on electromagnetic simulation tools to predict and avoid undesired modes in what will eventually become wafer-scale devices. Modeling these effects requires incorporating superconducting physics and the cryogenic properties of dielectrics into electromagnetic tools while developing methods to efficiently solve for high-Q resonances in large structures. Microwave expertise will be required, both in developing efficient simulation tools tailored to this purpose as well as in developing techniques to mitigate these undesired modes.

As the system size grows, it may also become important to develop techniques to predict system performance from within a single design environment (similar to the infrastructure that has been developed for digital design). This may involve developing qubit models which are compatible with commercial circuit solvers (e.g. [161]) and tying the classical and quantum systems together in a circuit solver tool. For systems such as superconducting quantum computers, which are spatially distributed, it may also be necessary to develop techniques to accurately model the interconnects between the classical control system and the quantum processor.

IX. CONCLUSION

Microwave technology has played a key role in the rise of quantum computing, and the two fields will continue to interact synergistically in the years ahead. The pursuit of an error-corrected quantum computer has been enabled by microwave technologies made available by the explosive growth of the wireless communication industry. However, reaching the performance currently achieved by today's state-of-the-art quantum processors has also necessitated the development of new microwave technologies, such as quantum-limited parametric amplifiers.

Today, there is a growing race to implement a fault-tolerant quantum computer. As described in this article, numerous microwave-related challenges must be overcome to build such a device, and continued microwave innovation will certainly be required. Critical areas for further research range from the development of new ultra-low-loss interconnect systems to the design of highly efficient quantum control systems. The

active engagement of microwave engineers in this exciting and important effort will be essential to its success.

ACKNOWLEDGMENT

The authors would like to thank J. Aumentado, K. Beloy, O. Naaman, H. Neven, and L. J. Stephenson for a careful reading of the manuscript, and M. Foss-Feig, D. Sank, and K. C. Young for helpful discussions.

REFERENCES

- [1] R. Buderl, *The Invention That Changed the World: How a Small Group of Radar Pioneers Won the Second World War and Launched a Technological Revolution* (Sloan Technology Series). New York, NY, USA: Simon & Schuster, 1996.
- [2] C. H. Townes and A. L. Schawlow, *Microwave Spectroscopy*. Mineola, NY, USA: Dover, 1975.
- [3] C. P. Slichter, *Principles of Magnetic Resonance* (Springer Series in Solid-State Sciences). Berlin, Germany: Springer, 1996.
- [4] E. M. Purcell, H. C. Torrey, and R. V. Pound, "Resonance absorption by nuclear magnetic moments in a solid," *Phys. Rev.*, vol. 69, pp. 37–38, 1946.
- [5] F. Bloch, W. W. Hansen, and M. Packard, "Nuclear induction," *Phys. Rev.*, vol. 69, pp. 127–127, 1946.
- [6] J. C. Bardin, D. Sank, O. Naaman, and E. Jeffrey, "Quantum computing: An introduction for microwave engineers," *IEEE Microw. Mag.*, vol. 21, no. 8, pp. 24–44, Aug. 2020.
- [7] F. Arute *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, U.K.: Cambridge Univ. Press, 2000.
- [9] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM Rev.*, vol. 41, no. 2, p. 303, 1999.
- [10] D. J. Wineland *et al.*, "Experimental issues in coherent quantum-state manipulation of trapped atomic ions," *J. Res. Nat. Inst. Stand. Technol.*, vol. 103, no. 3, pp. 259–328, 1998.
- [11] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, "Quantum dynamics of single trapped ions," *Rev. Mod. Phys.*, vol. 75, no. 1, pp. 281–324, 2003.
- [12] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, "Trapped-ion quantum computing: Progress and challenges," *Appl. Phys. Rev.*, vol. 6, no. 2, 2019, Art. no. 021314.
- [13] S. Seidelin *et al.*, "Microfabricated surface-electrode ion trap for scalable quantum information processing," *Phys. Rev. Lett.*, vol. 96, no. 25, 2006, Art. no. 253003.
- [14] K. Wright *et al.*, "Benchmarking an 11-qubit quantum computer," *Nature Commun.*, vol. 10, no. 1, 2019, Art. no. 5464.
- [15] J. M. Pino *et al.*, "Demonstration of the QCCD trapped-ion quantum computer architecture," 2020, *arXiv:2003.01293*.
- [16] J. Eschner, G. Morigi, F. Schmidt-Kaler, and R. Blatt, "Laser cooling of trapped ions," *J. Opt. Soc. Amer. B*, vol. 20, no. 5, pp. 1003–1015, 2003.
- [17] T. Ruster *et al.*, "A long-lived Zeeman trapped-ion qubit," *Appl. Phys. B*, vol. 122, no. 10, 2016, Art. no. 254.
- [18] T. P. Harty *et al.*, "High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit," *Phys. Rev. Lett.*, vol. 113, no. 22, 2014, Art. no. 220501.
- [19] P. Wang *et al.*, "Single ion-qubit exceeding one hour coherence time," 2020, *arXiv:2008.00251*.
- [20] H. Ball, W. D. Oliver, and M. J. Biercuk, "The role of master clock stability in quantum information processing," *npj Quantum Inf.*, vol. 2, no. 1, 2016, Art. no. 16033.
- [21] M. A. Sepiol *et al.*, "Probing qubit memory errors at the part-per-million level," *Phys. Rev. Lett.*, vol. 123, no. 11, 2019, Art. no. 110503.
- [22] B. E. Kane, "A silicon-based nuclear spin quantum computer," *Nature*, vol. 393, no. 6681, pp. 133–137, 1998.
- [23] F. A. Zwanenburg *et al.*, "Silicon quantum electronics," *Rev. Mod. Phys.*, vol. 85, pp. 961–1019, Jul. 2013.
- [24] A. Sigillito *et al.*, "Site-selective quantum control in an isotopically enriched $^{28}\text{Si}/\text{Si}_{0.7}\text{Ge}_{0.3}$ quadruple quantum dot," *Phys. Rev. Appl.*, vol. 11, Jun. 2019, Art. no. 061006.
- [25] J. Wrachtrup and F. Jelezko, "Processing quantum information in diamond," *J. Phys., Condens. Matter*, vol. 18, no. 21, pp. S807–S824, May 2006.
- [26] L. P. Kouwenhoven *et al.*, *Electron Transport Quantum Dots*. Dordrecht, The Netherlands: Springer Netherlands, 1997, pp. 105–214.
- [27] C. H. Yang *et al.*, "Operation of a silicon quantum processor unit cell above one kelvin," *Nature*, vol. 580, no. 7803, pp. 350–354, 2020.
- [28] L. Petit *et al.*, "Universal quantum logic in hot silicon qubits," *Nature*, vol. 580, no. 7803, pp. 355–359, 2020.
- [29] D. Loss and D. P. DiVincenzo, "Quantum computation with quantum dots," *Phys. Rev. A*, vol. 57, pp. 120–126, Jan. 1998.
- [30] A. Morello *et al.*, "Single-shot readout of an electron spin in silicon," *Nature*, vol. 467, no. 7316, pp. 687–91, 2010.
- [31] J. Levy, "Universal quantum computation with spin-1/2 pairs and Heisenberg exchange," *Phys. Rev. Lett.*, vol. 89, Sep. 2002, Art. no. 147902.
- [32] X. Croot *et al.*, "Device architecture for coupling spin qubits via an intermediate quantum state," *Phys. Rev. Appl.*, vol. 10, Oct. 2018, Art. no. 044058.
- [33] D. P. DiVincenzo, D. Bacon, J. Kempe, G. Burkard, and K. B. Whaley, "Universal quantum computation with the exchange interaction," *Nature*, vol. 408, no. 6810, pp. 339–342, 2000.
- [34] J. Medford *et al.*, "Quantum-dot-based resonant exchange qubit," *Phys. Rev. Lett.*, vol. 111, Jul. 2013, Art. no. 050501.
- [35] R. Hanson, L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M. K. Vandersypen, "Spins in few-electron quantum dots," *Rev. Mod. Phys.*, vol. 79, pp. 1217–1265, Oct. 2007.
- [36] J. M. Martinis, M. H. Devoret, and J. Clarke, "Energy-level quantization in the zero-voltage state of a current-biased Josephson junction," *Phys. Rev. Lett.*, vol. 55, no. 15, pp. 1543–1546, 1985.
- [37] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, "Coherent control of macroscopic quantum states in a single-Cooper-pair box," *Nature*, vol. 398, no. 6730, pp. 786–788, 1999.
- [38] J. Koch *et al.*, "Charge-insensitive qubit design derived from the Cooper pair box," *Phys. Rev. A*, vol. 76, no. 4, 2007, Art. no. 042319.
- [39] P. Krantz *et al.*, "A quantum engineer's guide to superconducting qubits," *Appl. Phys. Rev.*, vol. 6, no. 2, 2019, Art. no. 021318.
- [40] Z. Chen, "Metrology of quantum control and measurement in superconducting qubits," Ph.D. dissertation, UC Santa Barbara, 2018.
- [41] A. P. Place *et al.*, "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds," 2020, *arXiv:2003.00024*.
- [42] P. Jurcevic *et al.*, "Demonstration of quantum volume 64 on a superconducting quantum computing system," 2020, *arXiv:2008.08571*.
- [43] P. Klimov *et al.*, "Fluctuations of energy-relaxation times in superconducting qubits," *Phys. Rev. Lett.*, vol. 121, no. 9, 2018, Art. no. 090502.
- [44] F. H. L. Koppens *et al.*, "Driven coherent oscillations of a single electron spin in a quantum dot," *Nature*, vol. 442, no. 7104, pp. 766–771, 2006.
- [45] J. van Dijk *et al.*, "Impact of classical control electronics on qubit fidelity," *Phys. Rev. Appl.*, vol. 12, no. 4, 2019, Art. no. 044054.
- [46] C. Ospelkaus *et al.*, "Microwave quantum logic gates for trapped ions," *Nature*, vol. 476, no. 7359, pp. 181–184, 2011.
- [47] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, "Introduction to quantum noise, measurement, and amplification," *Rev. Modern Phys.*, vol. 82, no. 2, pp. 1155–1208, 2010.
- [48] T. A. Brun, "Quantum error correction," 2020, *arXiv:1910.03672*.
- [49] J. Roffe, "Quantum error correction: an introductory guide," *Contemporary Phys.*, vol. 60, no. 3, pp. 226–245, 2019.
- [50] E. Knill, "Quantum computing with realistically noisy devices," *Nature*, vol. 434, no. 7029, pp. 39–44, 2005.
- [51] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Phys. Rev. A*, vol. 86, no. 3, 2012, Art. no. 032324.
- [52] M. Steffen, J. M. Martinis, and I. L. Chuang, "Accurate control of Josephson phase qubits," *Phys. Rev. B*, vol. 68, no. 22, 2003, Art. no. 224518.
- [53] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, "Simple pulses for elimination of leakage in weakly nonlinear qubits," *Phys. Rev. Lett.*, vol. 103, no. 11, 2009, Art. no. 110501.

- [54] J. M. Chow *et al.*, “Optimized driving of superconducting artificial atoms for improved single-qubit gates,” *Phys. Rev. A*, vol. 82, no. 4, 2010, Art. no. 040305.
- [55] Z. Chen *et al.*, “Measuring and suppressing quantum state leakage in a superconducting qubit,” *Phys. Rev. Lett.*, vol. 116, no. 2, 2016, Art. no. 020501.
- [56] D. P. L. Aude Craik *et al.*, “High-fidelity spatial and polarization addressing of $^{43}\text{Ca}^+$ qubits using near-field microwave control,” *Phys. Rev. A*, vol. 95, no. 2, 2017, Art. no. 022337.
- [57] C. Piltz, T. Sriarunothai, A. F. Varón, and C. Wunderlich, “A trapped-ion-based quantum byte with 10^{-5} next-neighbour cross-talk,” *Nature Commun.*, vol. 5, no. 1, 2014, Art. no. 4679.
- [58] D. Leibfried, “Individual addressing and state readout of trapped ions utilizing RF micromotion,” *Phys. Rev. A*, vol. 60, no. 5, pp. R3335–R3338, 1999.
- [59] P. Staunum, and M. Drewsen, “Trapped-ion quantum logic utilizing position-dependent ac Stark shifts,” *Phys. Rev. A*, vol. 66, no. 4, 2002, Art. no. 040302.
- [60] U. Warring *et al.*, “Individual-ion addressing with microwave field gradients,” *Phys. Rev. Lett.*, vol. 110, no. 17, 2013, Art. no. 173002.
- [61] R. Srinivas, “Laser-free trapped-ion quantum logic with a radio-frequency magnetic field gradient,” Ph.D. dissertation, Univ. Colorado, Boulder, 2020.
- [62] K. Takeda, A. Noiri, J. Yoneda, T. Nakajima, and S. Tarucha, “Resonantly driven singlet-triplet spin qubit in silicon,” *Phys. Rev. Lett.*, vol. 124, Mar. 2020, Art. no. 117701.
- [63] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, “Efficient Z gates for quantum computing,” *Phys. Rev. A*, vol. 96, no. 2, 2017, Art. no. 022330.
- [64] R. Srinivas *et al.*, “Trapped-ion spin-motion coupling with microwaves and a near-motional oscillating magnetic field gradient,” *Phys. Rev. Lett.*, vol. 122, no. 16, 2019, Art. no. 163201.
- [65] R. Barends *et al.*, “Superconducting quantum circuits at the surface code threshold for fault tolerance,” *Nature*, vol. 508, no. 7497, pp. 500–503, 2014.
- [66] F. Yan *et al.*, “Tunable coupling scheme for implementing high-fidelity two-qubit gates,” *Phys. Rev. App.*, vol. 10, no. 5, 2018, Art. no. 054062.
- [67] B. Foxen *et al.*, “Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms,” 2020, *arXiv:2001.08343*.
- [68] C. Rigetti and M. Devoret, “Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies,” *Phys. Rev. B*, vol. 81, no. 13, 2010, Art. no. 134507.
- [69] M. D. Shulman *et al.*, “Demonstration of entanglement of electrostatically coupled singlet-triplet qubits,” *Science*, vol. 336, no. 6078, pp. 202–205, 2012.
- [70] A. A. Clerk, K. W. Lehnert, P. Bertet, J. R. Petta, and Y. Nakamura, “Hybrid quantum systems with circuit quantum electrodynamics,” *Nature Phys.*, vol. 16, no. 3, pp. 257–267, 2020.
- [71] F. Mintert and C. Wunderlich, “Ion-trap quantum logic using long-wavelength radiation,” *Phys. Rev. Lett.*, vol. 87, no. 25, 2001, Art. no. 257904.
- [72] M. Johanning *et al.*, “Individual addressing of trapped ions and coupling of motional and spin states using RF radiation,” *Phys. Rev. Lett.*, vol. 102, no. 7, 2009, Art. no. 073004.
- [73] R. T. Sutherland *et al.*, “Versatile laser-free trapped-ion entangling gates,” *New J. Phys.*, vol. 21, no. 3, 2019, Art. no. 033033.
- [74] C. Ospelkaus *et al.*, “Trapped-ion quantum logic gates based on oscillating magnetic fields,” *Phys. Rev. Lett.*, vol. 101, no. 9, 2008, Art. no. 090502.
- [75] U. Warring *et al.*, “Techniques for microwave near-field quantum control of trapped ions,” *Phys. Rev. A*, vol. 87, no. 1, 2013, Art. no. 013437.
- [76] J. Welzel, F. Stopp, and F. Schmidt-Kaler, “Spin and motion dynamics with zigzag ion crystals in transverse magnetic gradients,” *J. Phys. B, Atomic Mol. Opt. Phys.*, vol. 52, no. 2, 2019, Art. no. 025301.
- [77] T. P. Harty *et al.*, “High-fidelity trapped-ion quantum logic using near-field microwaves,” *Phys. Rev. Lett.*, vol. 117, no. 14, 2016, Art. no. 140501.
- [78] S. Weidt *et al.*, “Trapped-ion quantum logic with global radiation fields,” *Phys. Rev. Lett.*, vol. 117, no. 22, 2016, Art. no. 220501.
- [79] H. Hahn *et al.*, “Integrated $^9\text{Be}^+$ multi-qubit gate device for the ion-trap quantum computer,” *npj Quantum Inf.*, vol. 5, no. 1, 2019, Art. no. 70.
- [80] G. Zarantonello *et al.*, “Robust and resource-efficient microwave near-field entangling $^9\text{Be}^+$ gate,” *Phys. Rev. Lett.*, vol. 123, no. 26, 2019, Art. no. 260503.
- [81] V. M. Schäfer *et al.*, “Fast quantum logic gates with trapped-ion qubits,” *Nature*, vol. 555, no. 7694, pp. 75–78, 2018.
- [82] R. Ozeri *et al.*, “Errors in trapped-ion quantum gates due to spontaneous photon scattering,” *Phys. Rev. A*, vol. 75, no. 4, 2007, Art. no. 042329.
- [83] J. Rafferty *et al.*, “Direct digital synthesis of microwave waveforms for quantum computing,” 2017, *arXiv:1703.00942*.
- [84] W. D. Kalfus *et al.*, “High-fidelity control of superconducting qubits using direct microwave synthesis in higher Nyquist zones,” 2020, *arXiv:2008.02873*.
- [85] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, “Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation,” *Phys. Rev. A*, vol. 69, no. 6, 2004, Art. no. 062320.
- [86] A. Wallraff *et al.*, “Approaching unit visibility for control of a superconducting qubit with dispersive readout,” *Phys. Rev. Lett.*, vol. 95, no. 6, 2005, Art. no. 060501.
- [87] X. Mi, J. Cady, D. Zajac, P. Deelman, and J. Petta, “Strong coupling of a single electron in silicon to a microwave photon,” *Science*, vol. 355, no. 6321, pp. 156–158, 2017.
- [88] A. Stockklauser *et al.*, “Strong coupling cavity QED with gate-defined double quantum dots enabled by a high impedance resonator,” *Phys. Rev. X*, vol. 7, no. 1, 2017, Art. no. 011030.
- [89] A. West *et al.*, “Gate-based single-shot readout of spins in silicon,” *Nature Nanotechnol.*, vol. 14, no. 5, pp. 437–441, 2019.
- [90] M. Boissonneault, J. M. Gambetta, and A. Blais, “Dispersive regime of circuit QED: Photon-dependent qubit dephasing and relaxation rates,” *Phys. Rev. A*, vol. 79, no. 1, 2009, Art. no. 013819.
- [91] D. H. Slichter *et al.*, “Measurement-induced qubit state mixing in circuit QED from up-converted dephasing noise,” *Phys. Rev. Lett.*, vol. 109, no. 15, 2012, Art. no. 153601.
- [92] J. Aumentado, “Superconducting parametric amplifiers: The state of the art in Josephson parametric amplifiers,” *IEEE Microw. Mag.*, vol. 21, no. 8, pp. 45–59, Aug. 2020.
- [93] R. Vijay, D. H. Slichter, and I. Siddiqi, “Observation of quantum jumps in a superconducting artificial atom,” *Phys. Rev. Lett.*, vol. 106, no. 11, 2011, Art. no. 110502.
- [94] E. Jeffrey *et al.*, “Fast accurate state measurement with superconducting qubits,” *Phys. Rev. Lett.*, vol. 112, no. 19, 2014, Art. no. 190504.
- [95] M. D. Reed *et al.*, “Fast reset and suppressing spontaneous emission of a superconducting qubit,” *Appl. Phys. Lett.*, vol. 96, no. 20, 2010, Art. no. 203110.
- [96] A. Opremcak *et al.*, “Measurement of a superconducting qubit with a microwave photon counter,” *Science*, vol. 361, no. 6408, pp. 1239–1242, 2018.
- [97] A. Opremcak *et al.*, “High-fidelity measurement of a superconducting qubit using an on-chip microwave photon counter,” 2020, *arXiv:2008.02346*.
- [98] H. G. Dehmelt “Monoion oscillator as potential ultimate laser frequency standard,” *IEEE Trans. Instrum. Meas.*, vol. IM-31, no. 2, pp. 83–87, Jun. 1982.
- [99] A. H. Myerson *et al.*, “High-fidelity readout of trapped-ion qubits,” *Phys. Rev. Lett.*, vol. 100, no. 20, 2008, Art. no. 200502.
- [100] J. W. Britton *et al.*, “Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins,” *Nature*, vol. 484, no. 7395, pp. 489–492, 2012.
- [101] J. Dilling, K. Blaum, M. Brodeur, and S. Eliseev, “Penning-trap mass measurements in atomic and nuclear physics,” *Annu. Rev. Nucl. Particle Sci.*, vol. 68, no. 1, pp. 45–74, 2018.
- [102] S. Jain, J. Alonso, M. Grau, and J. P. Home, “Scalable arrays of micro-penning traps for quantum computing and simulation,” *Phys. Rev. X*, vol. 10, no. 3, 2020, Art. no. 031027.
- [103] W. Macalpine and R. Schildknecht, “Coaxial resonators with helical inner conductor,” *Proc. IRE*, vol. 47, no. 12, pp. 2099–2105, 1959.
- [104] J. D. Siverns, L. R. Simkins, S. Weidt, and W. K. Hensinger, “On the application of radio frequency voltages to ion traps via helical resonators,” *Appl. Phys. B*, vol. 107, no. 4, pp. 921–934, 2012.
- [105] S. R. Jefferts, C. Monroe, E. W. Bell, and D. J. Wineland, “Coaxial-resonator-driven RF (Paul) trap for strong confinement,” *Phys. Rev. A*, vol. 51, no. 4, pp. 3112–3116, 1995.

- [106] D. Gandolfi, M. Niedermayr, M. Kumph, M. Brownnutt, and R. Blatt, "Compact radio-frequency resonator for cryogenic ion traps," *Rev. Sci. Instrum.*, vol. 83, no. 8, 2012, Art. no. 084705.
- [107] M. F. Brandl, P. Schindler, T. Monz, and R. Blatt, "Cryogenic resonator design for trapped ion experiments in Paul traps," *Appl. Phys. B*, vol. 122, no. 6, 2016, Art. no. 157.
- [108] D. T. C. Allcock, "Surface-electrode ion traps for scalable quantum computing," Ph.D. dissertation, Oxford Univ., 2011.
- [109] T. P. Harty, "High-fidelity quantum logic in intermediate-field 43Ca^+ ," Ph.D. dissertation, Oxford Univ., 2013.
- [110] K. G. Johnson *et al.*, "Active stabilization of ion trap radiofrequency potentials," *Rev. Sci. Instrum.*, vol. 87, no. 5, 2016, Art. no. 053110.
- [111] M. Saffman, "Quantum computing with atomic qubits and Rydberg interactions: Progress and challenges," *J. Phys. B, Atomic Mol. Opt. Phys.*, vol. 49, no. 20, 2016, Art. no. 202001.
- [112] L. Henriot *et al.*, "Quantum computing with neutral atoms," *Quantum*, vol. 4, 2020, Art. no. 327.
- [113] V. V. Dobrovitski, G. D. Fuchs, A. L. Falk, C. Santori, and D. D. Awschalom, "Quantum control over single spins in diamond," *Annu. Rev. Condens. Matter Phys.*, vol. 4, no. 1, pp. 23–50, 2013.
- [114] M. W. Doherty *et al.*, "The nitrogen-vacancy colour centre in diamond," *Phys. Rep.*, vol. 528, no. 1, pp. 1–45, 2013.
- [115] A. Korpel, "Acousto-optics—A review of fundamentals," *Proc. IEEE*, vol. 69, no. 1, pp. 48–53, 1981.
- [116] L. Desmarais, *Applied Electro Optics*. Upper Saddle River, NJ, USA: Prentice-Hall, 1998.
- [117] R. V. Pound, "Electronic frequency stabilization of microwave oscillators," *Rev. Sci. Instrum.*, vol. 17, no. 11, pp. 490–505, 1946.
- [118] R. W. P. Drever *et al.*, "Laser phase and frequency stabilization using an optical resonator," *Appl. Phys. B, Photophys. Laser Chem.*, vol. 31, no. 2, pp. 97–105, 1983.
- [119] J. Mizrahi *et al.*, "Quantum control of qubits and atomic motion using ultrafast laser pulses," *Appl. Phys. B*, vol. 114, no. 1–2, pp. 45–61, 2014.
- [120] W. H. Louisell, *Coupled Mode and Parametric Electronics*. New York, NY, USA: Wiley, 1960.
- [121] W. W. Mumford, "Some notes on the history of parametric transducers," *Proc. IRE*, vol. 48, no. 5, pp. 848–853, 1960.
- [122] M. A. Castellanos-Beltran and K. W. Lehnert, "Widely tunable parametric amplifier based on a superconducting quantum interference device array resonator," *Appl. Phys. Lett.*, vol. 91, no. 8, 2007, Art. no. 083509.
- [123] T. Yamamoto *et al.*, "Flux-driven Josephson parametric amplifier," *Appl. Phys. Lett.*, vol. 93, no. 4, 2008, Art. no. 042510.
- [124] R. Vijay, "Josephson bifurcation amplifier: Amplifying quantum signals using a dynamical bifurcation," Ph.D. dissertation, Yale Univ., 2008.
- [125] N. Bergeal *et al.*, "Phase-preserving amplification near the quantum limit with a Josephson ring modulator," *Nature*, vol. 465, no. 7294, pp. 64–68, 2010.
- [126] M. Hatridge *et al.*, "Quantum back-action of an individual variable-strength measurement," *Science*, vol. 339, no. 6116, pp. 178–181, 2013.
- [127] J. Y. Mutus *et al.*, "Strong environmental coupling in a Josephson parametric amplifier," *Appl. Phys. Lett.*, vol. 104, no. 26, 2014, Art. no. 263513.
- [128] T. Roy *et al.*, "Broadband parametric amplification with impedance engineering: Beyond the gain-bandwidth product," *Appl. Phys. Lett.*, vol. 107, no. 26, 2015, Art. no. 262601.
- [129] O. Naaman *et al.*, "High saturation power Josephson parametric amplifier with GHz bandwidth," in *Proc. IEEE MTT-S Int. Microw. Symp.*, vol. 2019, Jun. 2019, pp. 259–262.
- [130] C. Macklin *et al.*, "A near-quantum-limited Josephson traveling-wave parametric amplifier," *Science*, vol. 350, no. 6258, pp. 307–310, 2015.
- [131] B. Abdo, K. Sliwa, L. Frunzio, and M. Devoret, "Directional amplification with a Josephson circuit," *Phys. Rev. X*, vol. 3, Feb. 2013, Art. no. 031001.
- [132] F. Lecocq *et al.*, "Nonreciprocal microwave signal processing with a field-programmable Josephson amplifier," *Phys. Rev. Appl.*, vol. 7, no. 2, 2017, Art. no. 024028.
- [133] L. Ranzani and J. Aumentado, "Circulators at the quantum limit: Recent realizations of quantum-limited superconducting circulators and related approaches," *IEEE Microw. Mag.*, vol. 20, no. 4, pp. 112–122, Apr. 2019.
- [134] A. C. Mahoney *et al.*, "On-chip microwave quantum hall circulator," *Phys. Rev. X*, vol. 7, no. 1, 2017, Art. no. 011007.
- [135] K. M. Sliwa *et al.*, "Reconfigurable Josephson circulator / directional amplifier," *Phys. Rev. X*, vol. 5, Aug. 2015, Art. no. 041020.
- [136] F. Lecocq *et al.*, "Microwave measurement beyond the quantum limit with a nonreciprocal amplifier," *Phys. Rev. Appl.*, vol. 13, no. 4, 2020, Art. no. 044005.
- [137] B. Abdo, O. Jinka, N. T. Bronn, S. Olivadese, and M. Brink, "On-chip single-pump interferometric Josephson isolator for quantum measurements," 2020, *arXiv:2006.01918*.
- [138] E. I. Rosenthal *et al.*, "Efficient and low-backaction quantum measurement using a chip-scale detector," 2020, *arXiv:2008.03805*.
- [139] F. Lecocq *et al.*, "Efficient qubit measurement with a nonreciprocal microwave amplifier," 2020, *arXiv:2009.08863*.
- [140] S. Bosco, F. Haupt, and D. P. DiVincenzo, "Self-impedance-matched hall-effect gyrators and circulators," *Phys. Rev. Appl.*, vol. 7, Feb. 2017, Art. no. 024030.
- [141] A. C. Mahoney *et al.*, "Zero-field edge plasmons in a magnetic topological insulator," *Nature Commun.*, vol. 8, no. 1, 2017, Art. no. 1836.
- [142] S. Bosco, D. DiVincenzo, and D. Reilly, "Transmission lines and metamaterials based on quantum hall plasmonics," *Phys. Rev. Appl.*, vol. 12, Jul. 2019, Art. no. 014030.
- [143] Q. P. Herr, A. D. Smith, and M. S. Wire, "High speed data link between digital superconductor chips," *Appl. Phys. Lett.*, vol. 80, no. 17, pp. 3210–3212, 2002.
- [144] D. P. Franke, J. S. Clarke, L. M. Vandersypen, and M. Veldhorst, "Rents rule and extensibility in quantum computing," *Microprocessors Microsyst.*, vol. 67, pp. 1–7, 2019.
- [145] J. Hornibrook *et al.*, "Cryogenic control architecture for large-scale quantum computing," *Phys. Rev. Appl.*, vol. 3, no. 2, 2015, Art. no. 024010.
- [146] S. Pauka *et al.*, "Characterizing quantum devices at scale with custom cryo-CMOS," *Phys. Rev. Appl.*, vol. 13, no. 5, 2020, Art. no. 054072.
- [147] D. Reilly "Challenges in scaling-up the control interface of a quantum computer," in *Proc. IEEE Int. Electron Devices Meet.*, 2019, pp. 31–7.
- [148] S. Pauka *et al.*, "A cryogenic interface for controlling many qubits," 2019, *arXiv:1912.01299*.
- [149] J. Stuart *et al.*, "Chip-integrated voltage sources for control of trapped ions," *Phys. Rev. Appl.*, vol. 11, no. 2, 2019, Art. no. 024010.
- [150] B. Patra *et al.*, "Cryo-CMOS circuits and systems for quantum computing applications," *IEEE J. Solid-State Circuits*, vol. 53, no. 1, pp. 309–321, Jan. 2017.
- [151] B. Patra *et al.*, "A scalable cryo-CMOS 2-to-20 GHz digitally-intensive controller for 4×32 frequency multiplexed spin qubits/transmons in 22-nm FinFET technology for quantum computers," in *Proc. Int. Solid-State Circuits Conf.*, 2020, pp. 304–306.
- [152] J. C. Bardin *et al.*, "A 28 nm bulk-CMOS 4-to-8 GHz 2 mW cryogenic pulse modulator for scalable quantum computing," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2019, pp. 456–458.
- [153] J. C. Bardin *et al.*, "Design and characterization of a 28-nm bulk-cmos cryogenic quantum controller dissipating less than 2 mW at 3 K," *IEEE J. Solid-State Circuits*, vol. 54, no. 11, pp. 3043–3060, 2019.
- [154] I. Bashir *et al.*, "A mixed-signal control core for a fully integrated semiconductor quantum computer system-on-chip," in *Proc. ESSCIRC IEEE 45th Eur. Solid State Circuits Conf.*, 2019, pp. 125–128.
- [155] C. Degenhardt *et al.*, "Systems engineering of cryogenic CMOS electronics for scalable quantum computers," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2019, pp. 1–5.
- [156] N. Langford *et al.*, "Experimentally simulating the dynamics of quantum light and matter at deep-strong coupling," *Nature Commun.*, vol. 8, no. 1, 2017, Art. no. 1715.
- [157] M. Rol *et al.*, "Fast, high-fidelity conditional-phase gate exploiting leakage interference in weakly anharmonic superconducting qubits," *Phys. Rev. Lett.*, vol. 123, no. 12, 2019, Art. no. 120502.
- [158] *LNF-LNC4_8C 4–8 GHz Cryogenic Low Noise Amplifier*, Low Noise Factory, Göteborg, Sweden. [Online]. Available: https://www.lownoisefactory.com/files/7015/7825/6000/LNF-LNC4_8C.pdf
- [159] S. Montazeri, W.-T. Wong, A. H. Coskun, and J. C. Bardin, "Ultra-low-power cryogenic SiGe low-noise amplifiers: Theory and demonstration," *IEEE Trans. Microw. Theory Techn.*, vol. 64, no. 1, pp. 178–187, Jan. 2016.


- [160] W.-T. Wong, M. Hosseini, H. Rucker, and J. C. Bardin, "A 1 mW cryogenic LNA exploiting optimized SiGe HBTs to achieve an average noise temperature of 3.2 K from 4-8 GHz," in *Proc. IEEE IMS*, 2020, pp. 181–184.
- [161] J. v. Dijk, A. Vladimirescu, M. Babaie, E. Charbon, and F. Sebastiano, "Spine (spin emulator) - A quantum-electronics interface simulator," in *Proc. IEEE 8th Int. Workshop Adv. Sensors Interfaces*, 2019, pp. 23–28.



JOSEPH C. BARDIN (Senior Member, IEEE) received the B.S. degree from the University of California, Santa Barbara, in 2003, the M.S. degree from the University of California, Los Angeles, in 2005, and the Ph.D. degree from the California Institute of Technology, in 2009, all in electrical engineering.

From 2003 to 2005, he was with the Jet Propulsion Laboratory, California Institute of Technology, where he was involved in the demonstration of an array-based downlink for the NASA deep-space network. From 2009 to 2010, he was a Postdoctoral Researcher with the Caltech High-Speed Integrated Circuits Group, where he focused on selfhealing microwave receivers. In 2010, he joined the University of Massachusetts Amherst, where he is currently a Professor of Electrical and Computer Engineering. Since 2017, he has also been with Google Quantum AI, where he focuses on the development of novel integrated electronics for quantum computing.

Dr. Bardin currently is an Associate Editor and Steering Committee Member for IEEE TRANSACTIONS ON QUANTUM ENGINEERING and a Track Editor for IEEE JOURNAL OF MICROWAVES. He was the recipient of the 2011 DARPA Young Faculty Award, the 2014 NSF CAREER Award, the 2015 ONR YIP Award, the 2016 UMass College of Engineering Outstanding Junior Faculty Award, the 2016 UMass Convocation Award for Excellence in Research and Creative Activities, and the 2020 IEEE MTT-S Outstanding Young Engineer Award.



DANIEL H. SLICHTER (Senior Member, IEEE) received the A.B. (*magna cum laude*) degree from Harvard University in 2004, the M.A. degree from the University of California, Berkeley, in 2007, and the Ph.D. degree from the University of California, in 2011, all in physics.

From 2007 to 2011, he was a researcher with the University of California, Berkeley, Quantum Nanoelectronics Laboratory, where he demonstrated the first continuous high-fidelity state readout of a superconducting qubit by using a near-quantum-

limited microwave-frequency superconducting parametric amplifier, and studied quantum measurement backaction and quantum feedback. Since 2012, he has been with the National Institute of Standards and Technology (NIST), Boulder, CO, first as a Postdoctoral Research Fellow and currently as a Staff Physicist, where his research has focused on high-fidelity microwave-based control and entanglement of trapped ion qubits, as well as trap-integrated superconducting photon detectors for trapped ion qubit readout.

Dr. Slichter was the recipient of a Hertz Foundation Fellowship (2006–2011), the Hertz Foundation Thesis Prize (2012), an NRC Postdoctoral Fellowship (2012–2014), and a Kavli Fellowship (2016). He currently is an Associate Editor for IEEE TRANSACTIONS ON QUANTUM ENGINEERING.



DAVID J. REILLY (Member, IEEE) received the B. App. Sci. (H1) degree in applied physics from the University of Technology, Sydney (UTS) in 1998. From 1995 to 1998, he held a Technical Staff Position with CSIRO (Applied Physics), characterising high T_c superconducting films for RF and microwave applications. In 2003, he was awarded the Ph.D. degree in physics from the University of New South Wales (UNSW), for work related to quantum transport in mesoscale semiconductor devices.

From 2003 to 2005, he was a Hewlett-Packard Fellow in quantum computing with UNSW, employing high frequency charge sensing measurements for the readout of superconducting qubits and silicon spin qubits. From 2005 to 2008, he continued this research theme at Harvard University, developing microwave approaches for detecting the state of spin qubits in a single shot as well as probing spin dynamics. Since 2008, he has been a Faculty member of the School of Physics at the University of Sydney, where he is currently Professor. Dr. Reilly is also a Chief Investigator in the ARC Centre of Excellence for Engineered Quantum Systems (EQUS), where his research focuses on developing quantum technology that spans computing, sensing, and new biomedical imaging modalities. In 2017, he joined Microsoft Corp. where he is a Partner Research Manager and Director of Microsoft Quantum Sydney. As a leader in Microsoft's efforts in quantum computing, much of his work is at the quantum-classical interface and the scale-up of quantum technology.