

Simulating nonlinear optical processes on a superconducting quantum device

Yuan Shi^{1,†}, Bram Evert², Amy F. Brown³, Vinay Tripathi³, Eyob A. Sete², Vasily Geyko⁴, Yujin Cho⁴, Jonathan L. DuBois⁴, Daniel Lidar^{3,5}, Ilon Joseph⁴ and Matt Reagor²

¹Department of Physics, Center for Integrated Plasma Studies, University of Colorado Boulder, Boulder, CO 80309, USA

²Rigetti Computing, 775 Heinz Avenue, Berkeley, CA 94710, USA

³Department of Physics and Astronomy and Center for Quantum Information Science and Technology, University of Southern California, Los Angeles, CA 90089, USA

⁴Lawrence Livermore National Laboratory, Livermore, CA 94550, USA

⁵Department of Electrical and Computer Engineering and Department of Chemistry, University of Southern California, Los Angeles, CA 90089, USA

(Received 18 June 2024; revised 20 September 2024; accepted 27 September 2024)

Simulating plasma physics on quantum computers is difficult because most problems of interest are nonlinear, but quantum computers are not naturally suitable for nonlinear operations. In weakly nonlinear regimes, plasma problems can be modelled as wave–wave interactions. In this paper, we develop a quantization approach to convert nonlinear wave–wave interaction problems to Hamiltonian simulation problems. We demonstrate our approach using two qubits on a superconducting device. Unlike a photonic device, a superconducting device does not naturally have the desired interactions in its native Hamiltonian. Nevertheless, Hamiltonian simulations can still be performed by decomposing required unitary operations into native gates. To improve experimental results, we employ a range of error-mitigation techniques. Apart from readout error mitigation, we use randomized compilation to transform undiagnosed coherent errors into well-behaved stochastic Pauli channels. Moreover, to compensate for stochastic noise, we rescale exponentially decaying probability amplitudes using rates measured from cycle benchmarking. We carefully consider how different choices of product-formula algorithms affect the overall error and show how a trade-off can be made to best utilize limited quantum resources. This study provides an example of how plasma problems may be solved on near-term quantum computing platforms.

Key words: plasma nonlinear phenomena, plasma simulation, quantum plasma

† Email address for correspondence: yuan.shi@colorado.edu

1. Introduction

The physics of plasmas is often multi-scale, multi-physics and highly nonlinear. While classical computers have no trouble handling nonlinearity, the multi-scale and multi-physics aspects make many plasma problems challenging even for classical supercomputers. Simulations of larger systems with finer resolutions are restricted by both the memory and time constraints of classical computers. In comparison, quantum computers, which are still under development, in principle have an exponentially larger memory, and thus have attracted substantial interest in recent years, including within the plasma community (Dodin & Startsev 2021; Joseph *et al.* 2023; Koukoutsis *et al.* 2023). In addition to a larger memory, quantum computers support algorithms that offer quadratic to exponential speedups over the best-known classical algorithms for specialized problems, such as quantum search, Fourier transform and factoring. However, whether quantum computers have the potential to offer a speedup for plasma problems remains an open question. A major difficulty is that quantum algorithms typically rely on unitary operations, and quantum computers are not naturally suited for general nonlinear operations needed in plasma physics. Attempts have been made to develop schemes that can handle nonlinear plasma problems on quantum computers (Joseph 2020; Engel, Smith & Parker 2021; Liu *et al.* 2021; Lin *et al.* 2022). Many of these schemes rely on future fault-tolerant quantum computers and thus cannot be tested on current noisy quantum devices. While developing abstract algorithms for future quantum hardware is valuable, performing concrete examples on current devices allows the community to build intuition about how quantum computation works in practice and reveals potential gaps between theoretical expectations and the reality of present-day capabilities.

In this paper, we show that small initial value problems that model the evolution of nonlinear wave–wave interactions can be solved on current quantum hardware, after employing a suite of error-mitigation techniques. Wave–wave interactions are a general framework for weakly nonlinear dynamical systems (Zhakarov, L’vov & Falkovich 1992; Nazarenko 2011; Davidson 2012). Many dynamical systems possess fixed points, which correspond to equilibrium states of the system. When slightly perturbed away from a stable fixed point, the system oscillates. Small-amplitude oscillations of fields are usually called linear waves in plasma physics. At the linear level, the waves are decoupled eigenmodes. However, at larger amplitudes, waves become coupled due to nonlinearities of the system, leading to what are known as wave–wave interactions. For example, laser–plasma interactions, such as Raman and Brillouin scattering, are often described as three-wave interactions (Michel 2023), because they involve an incoming laser, an outgoing laser and a mediating plasma wave. As another example, filamentation and modulational instabilities of plasma waves can be described as four-wave interactions (Michel 2023). In these instabilities, fluctuations of the laser amplitude depend on the laser intensity. The instability is a four-wave process because the laser beats with its fluctuations to produce side bands. Wave–wave interactions not only arise in laser–plasma interactions, but also occur in magnetically confined plasmas (Nazarenko 2011; Hansen *et al.* 2017), astrophysical plasmas (Bowen *et al.* 2018) and many other physical systems (Zhakarov *et al.* 1992; Nazarenko 2011).

While plasma physics often treats wave–wave interactions classically, we note that these nonlinear processes are intrinsically quantum. In fact, Raman (Raman & Krishnan 1928) and Brillouin (Brillouin 1914) scattering were first studied in gases and solids as quantum processes before these terms were borrowed by plasma physicists. For example, at the quantum level, a single incoming photon can scatter into a single outgoing photon by emitting a single phonon (Bowen & Milburn 2015). When there are many indistinguishable

photons, the scattering probability amplitudes add up to determine the total scattering probability, but the underlying three-wave interaction is not fundamentally different. In the study of laser–plasma interactions, laser light is often treated as classical electromagnetic waves. However, it is worth noting that from a quantum perspective, laser light is just a special collection of photons known as a coherent state, for which photon number distribution is Poissonian. In this regard, laser scattering from plasma is just a special example of three-wave interactions. More generally, photons can occupy other quantum states (Loudon 2000). For example, other important classes of quantum states are known as squeezed states (Breitenbach, Schiller & Mlynek 1997), which are different from coherent states because their probability distribution functions are not Poissonian. While laser light may be approximated as classical electromagnetic waves, squeezed light cannot be described by classical waveforms at all. If one attempts to assign a waveform to a squeezed state, then the waveform would have to fluctuate from cycle to cycle, and yet its Fourier transform must still have a narrow peak. To model quantum light interacting with matter, one can no longer use a classical treatment of wave–wave interactions. To capture quantum interference, which results from summing probability amplitudes rather than the probabilities themselves, one must use a quantum treatment.

The high cost associated with modelling the exact quantum mechanical evolution makes it desirable to develop quantum simulation capabilities for wave–wave interactions induced by quantum light. At the same time, the intrinsic quantum nature of wave–wave interactions makes it natural to study these nonlinear processes using a quantization approach (Shi, Qin & Fisch 2017; Shi *et al.* 2021a; Shi, Qin & Fisch 2021b). After promoting classical amplitudes to quantum creation and annihilation operators, the nonlinear equations that are used to describe wave–wave interactions are derivable as Heisenberg equations. By choosing a convenient basis, the Hamiltonian operators can be converted to finite-dimensional Hamiltonian matrices, and we perform quantum Hamiltonian simulations in the Schrödinger picture. The goal of the simulations is to extract observables at some later time, which can be constructed using the quantum state we evolve. In the limit of large occupation number, the quantized model reproduces linear instabilities in classical plasmas (May & Qin 2023).

For quantum Hamiltonian simulations to be efficient, a number of criteria must be met. One criterion is that the Hamiltonian matrices must have special properties. In our case, the matrices are sparse, for which efficient quantum algorithms exist (Berry *et al.* 2007, 2014; Low & Chuang 2017), at least for future error-corrected quantum computers. Another criterion is that one must not be interested in the full information stored in the quantum memory. In our cases, we are interested, for example, in the intensity of the backscattered light, which involves just a few expectation values. Overall, the scheme we develop is efficient for simulating how quantum light interacts with plasmas and can accelerate such simulations on future error-corrected quantum computers.

We demonstrate our scheme on a current superconducting quantum device using a two-qubit example. Our demonstration pushes the limit of what current hardware can do and addresses important trade-off between different error sources for product-formula algorithms. Rather than implementing black-box quantum Hamiltonian simulation algorithms, which are not feasible on current devices and will only become more efficient for larger system sizes, we simply decompose the requisite unitary operations into elementary gates using a Cartan decomposition (Smith *et al.* 2020; Huang *et al.* 2023), which is the most efficient approach for two-qubit problems. For more qubits, the number of two-qubit gates that Cartan decomposition requires grows exponentially, so implementing larger problems on future devices requires different methods. In our previous work, we have demonstrated that current devices can perform two-qubit time

evolution using the exact unitary (Shi *et al.* 2021a), which is obtained by analytically exponentiating the full Hamiltonian matrix. In principle, if one could perform the exact exponentiation, then the problem is already solved. However, even when the exact unitary was prescribed, we found that quantum devices had difficulties repeating the dense unitary beyond a few cycles, unless one can drastically improve gate fidelity.

In this paper, we demonstrate the next level of quantum computation by assuming that only certain parts of the exact unitary can be performed efficiently. More specifically, since our Hamiltonian has two non-commuting terms, we assume that each term can be exponentiated exactly. Then, we use product formulas to approximate the total unitary. Product formulas are often known as Hamiltonian splitting in the plasma literature (He *et al.* 2015; Morrison 2017), and are known as Lie–Trotter–Suzuki decomposition in the quantum literature (Trotter 1959; Suzuki 1976). On future quantum computers, quantum Hamiltonian simulations may still utilize product formulas, when the full Hamiltonian cannot be simulated efficiently but its subparts can.

For current devices, which do not yet support operational error correction, we show that reasonable results can be obtained only after we employ a range of error-mitigation and error-suppression techniques. First, we suppress the occurrence of errors by addressing the highest performing two-qubit gates on the device. Rather than using an off-the-shelf gate, we calibrate a \sqrt{i} SWAP (SQISW) gate (Abrams *et al.* 2020) which provides superior performance and expressiveness (Peterson, Crooks & Smith 2020) for our problem. Second, we mitigate readout errors, which occur if the qubit states are misclassified during measurements. We experimentally measure the confusion matrix, which describes the probability of preparing a qubit register in one state but classifying it in another state. We estimate the true distribution of samples from a noisy one using iterative Bayesian unfolding (Nachman *et al.* 2020). Third, we mitigate coherent gate errors, which occur if the realized unitary gate systematically differs from the target unitary. Coherent errors often arise from drifts in system parameters, so that gates calibrated at one time no longer remain perfectly calibrated at a later time. Coherent errors are problematic because they accumulate at each gate operation and can interfere constructively. To mitigate coherent errors, we use randomized compilation (Wallman & Emerson 2016; Hashim *et al.* 2020). When a unitary operation is called, the hardware implements it using a different but equivalent gate sequence, sampled at random from a precompiled set of choices. This technique converts coherent errors into stochastic errors, thereby suppressing constructive interference between errors. Finally, to mitigate incoherent errors, which occur due to both the random selection of gate sequences and intrinsic quantum decoherence on the hardware, we employ an amplitude-rescaling technique (Ville *et al.* 2022). The technique assumes that the probability of a pure quantum state decays exponentially to the fully mixed state. We measure the decay rate using cycle benchmarking (Erhard *et al.* 2019), and compensate for the decay by multiplying the probabilities with an exponential growth factor. While this technique partially recovers signals, it also leads to an exponential growth of error bars. The maximum simulation time and the maximum achievable gate depth are reached when the error bars become comparable to the signal of interest.

Using error-mitigation techniques, we substantially extend the achievable simulation depth, which is nevertheless still limited. After error mitigation, we can accurately measure observables at a depth of about two hundred two-qubit gates, for a reasonable shot budget. Two-qubit gates, which require longer hardware runtime and are more vulnerable to decoherence, have significantly lower fidelity than single-qubit gates, and are the limiting factor on current quantum devices. Given that errors grow with the gate depth, more accurate results may be obtained for a targeted final time by reducing the time step size, which in turn reduces the discretization error per step at the expense of increasing the

accumulation of hardware errors. Alternatively, more accurate results may be obtained by using a higher-order algorithm, which decreases the algorithmic error per step at the expense of more gates per step, but allows the use of larger time step sizes. We carefully consider different choices and show that a trade-off can be made to best utilize limited quantum resources.

The paper is organized as follows. In § 2, we discuss classical models of laser–plasma interactions and show how the model can be quantized and converted to a Hamiltonian simulation problem for quantum computers. In § 3, we implement an exact two-qubit problem on superconducting hardware, and show that error-mitigation techniques meaningfully improve simulation results. In § 4, we investigate product-formula approximations and discuss how to best utilize limited quantum resources.

2. Classical and quantum models of laser pulse compression

An important class of wave–wave interactions in plasma physics are laser–plasma interactions. In this paper, we consider an example scenario where a plasma is used for laser pulse compression (Malkin, Shvets & Fisch 1999), during which the intensity of a seed laser pulse is amplified while its duration is shortened. Other laser–plasma interaction scenarios, as well as other cases of wave–wave interactions, may be treated in a similar fashion. We first describe how the problem is usually treated classically in plasma physics, and then develop a quantized model that is amenable to simulations on quantum computers.

2.1. Classical model

Laser amplification is often treated as a parametric process, where the signal and idler waves grow by consuming a pump wave. When the pump energy is being replenished, or when the pump energy dominates, one may approximate the pump amplitude a_1 as a constant, in which case the seed amplitude a_2 and the idler amplitude a_3 grow exponentially. However, when the pump amplitude is not held constant, the three-wave nature of the underlying interaction becomes apparent. The interaction is often described by the three-wave equations (Davidson 2012)

$$d_t a_1 = g a_2 a_3, \tag{2.1a}$$

$$d_t a_2 = -g^* a_1 a_3^\dagger, \tag{2.1b}$$

$$d_t a_3 = -g^* a_1 a_2^\dagger, \tag{2.1c}$$

where d_t is the advective derivative, g is the coupling coefficient, g^* is its complex conjugate and a^\dagger denotes the complex conjugate of a in the classical model. The advective derivative is specific for each wave and is defined as $d_t = \partial_t + \mathbf{v}_g \cdot \nabla + \mu$, where $\mathbf{v}_g = \partial\omega/\partial\mathbf{k}$ is the group velocity and μ is the damping rate of the wave. The complex-valued amplitude a is the slowly varying envelope of the classical wave. The amplitude is normalized such that $n = |a|^2$ is proportional to the wave action density, which is proportional to the number of photons in the wave.

Three-wave interactions satisfy a number of conservation laws. First, the equations describe resonant interactions where both energy and momentum are conserved. For waves with narrow bandwidth, their 4-momentum density is proportional to $k^\mu = (\omega/c, \mathbf{k})$. Energy and momentum are conserved because the interaction satisfies the resonance conditions $k_1^\mu = k_2^\mu + k_3^\mu$. Moreover, the interaction has two independently conserved actions $S_2 = n_1 + n_2$ and $S_3 = n_1 + n_3$. These actions are constants of motion because their advective derivatives are zero. The physical meaning of action conservation is that

the three-wave interaction splits each pump photon into a seed and an idler photon. So, whenever n_1 decreases by one, both n_2 and n_3 increase by one, and vice versa.

The three-wave equations are partial differential equations that describe how the wave amplitudes evolve in both space and time. If all amplitudes are initially uniform in space, then they will remain uniform as they evolve in time. However, if the amplitudes are not uniform, two competing effects change the envelopes of the waves. First, wave advection transports the wave envelope in the direction of the group velocity. The advection is a linear process, and the envelope remains unchanged in the co-moving frame. Second, the three-wave interaction changes amplitudes locally. Since the interaction is nonlinear, the change is faster where the amplitudes of the other two waves are larger.

Laser pulse compression is a special scenario where the competition between the two effects leads to the amplification and shortening of the seed wave. At later stages of pulse compression, the intensity of the seed far exceeds that of the pump. The large a_2 induces an additional relativistic nonlinearity. The nonlinearity originates from the fact that, in plasmas, photons are massive particles due to their interactions with free charges. In unmagnetized plasmas, photons satisfy the dispersion relation $\omega^2 = \omega_p^2 + c^2k^2$, where the photon mass can be identified with the plasma frequency $\omega_p = (e^2n_e/\epsilon_0m_e)^{1/2}$. Here, e is the electron charge and n_e is the electron density. Because electrons oscillate in the laser's electric field, the effective electron mass m_e is replaced by γm_e when the electron quiver speed v_q becomes comparable to the speed of light c , where $\gamma = 1/\sqrt{1 - v_q^2/c^2}$.

As the seed pulse propagates, at places where the pulse is more intense, the photon mass $\omega_p \propto \gamma^{-1/2}$ becomes smaller. A smaller ω_p leads to a larger k at a fixed ω , which means a larger group velocity $v_g = c^2k/\omega$. Consequently, the more intense part of a_2 moves at a higher group velocity. If the envelope of a_2 has initial modulations, then they will pile up and grow. This process is known as relativistic modulational instability.

The modulational instability may be understood as a four-wave process, where the laser beats with its modulations to produce side bands. In the weakly relativistic limit, we can expand $1/\gamma \simeq 1 - v_q^2/(2c^2)$. Since electrons respond primarily to the electric field of the laser pulse, Newton's equation $m_e dv_q/dt \simeq eE$ becomes $\tilde{v}_q \simeq ie\tilde{E}/(m_e\omega)$ in Fourier space. Denoting the normalized laser amplitude by $\alpha = e\tilde{E}/(m_e\omega c)$, then the average quiver speed is $\langle v_q^2/c^2 \rangle = \frac{1}{2}e^2|\tilde{E}|^2/(m_e\omega c)^2 = \frac{1}{2}|\alpha|^2$. Replacing $\omega_p^2 \rightarrow \omega_p^2/\gamma \simeq \omega_p^2(1 - \frac{1}{4}|\alpha|^2)$, the photon dispersion relation is approximated. The dispersion relation is derivable from the wave equation $[\partial_t^2 - c^2\nabla^2 + \omega_p^2(1 - \frac{1}{4}|\alpha|^2)]E = 0$. In the Wentzel–Kramers–Brillouin approximation, the complex wave is $E(x, t) = A(x, t)e^{i\theta}$, where $\theta = k \cdot x - \omega t$, and the envelope A varies slowly in the sense that $|\partial_t A/A| \ll \omega$ and $|\nabla A/A| \ll k$. Then, to leading order, the wave equation is approximated by $[\partial_t + \mathbf{v}_g \cdot \nabla - i\omega_p^2|\alpha|^2/(8\omega)]A = 0$. When focusing on the scalar amplitude A , because $A \propto \alpha \propto a$, the equation that describes the modulational instability of the seed pulse is

$$d_t a_2 = iR a_2^\dagger a_2 a_2, \quad (2.2)$$

where d_t is again the advective derivative and $R = \omega_p^2/(8\omega)$ is the coupling coefficient. For relativistic modulational instability, $R > 0$ is a real number, which means $iR|a_2|^2$ is purely imaginary. The above equation thus modulates the phase of the complex a_2 in such a way that a larger $|a_2|$ leads to a faster phase evolution.

Laser pulse compression is described by the combined equations (2.1) and (2.2). The modulational instability only affects a_2 because during laser pulse compression, the peak values satisfy $|a_2| \gg |a_1|, |a_3|$. Since the three-wave interaction is a phase-sensitive

process, the modulational instability spoils the amplification process by introducing a phase mismatch.

2.2. Quantum model

In the classical model, the amplitude a is a complex-valued function, and $n = |a|^2$ is proportional to the number of photons. This set-up naturally admits canonical quantization for bosonic quantum fields $[a_i(\mathbf{x}), a_j^\dagger(\mathbf{y})] = \delta_{ij}(2\pi)^3\delta^{(3)}(\mathbf{x} - \mathbf{y})$, where the indices $i, j = 1, 2, 3$, and the operators have spatial dependencies. Since we will later implement the model on quantum hardware, which has a limited number of qubits, in this paper we focus on the temporal problem with no spatial dependence. In this case, when damping is negligible, the advective derivative $d_t \rightarrow \partial_t$ is reduced to a partial derivative in time, and the operators satisfy the canonical quantization conditions

$$[a_i, a_j^\dagger] = \delta_{ij}. \tag{2.3}$$

The quantization promotes normalized amplitudes to creation and annihilation operators, and the Kronecker delta distinguishes the three types of waves in the system. For each wave type, the number operator is $n_i = a_i^\dagger a_i$. The eigenstates of n_i are the Fock states $|m_i\rangle$, namely $n_i|m_i\rangle = m_i|m_i\rangle$, where $m_i = 0, 1, 2, \dots$ and $|m_i\rangle = (a_i^\dagger)^{m_i}|0_i\rangle/\sqrt{m_i!}$. Here, $|0_i\rangle$ is the ground state of wave i , which is annihilated by $a_i|0_i\rangle = 0$. More generally, these quantum harmonic oscillators have the usual matrix elements $a_i|m_i\rangle = \sqrt{m_i}|m_i - 1\rangle$ and $a_i^\dagger|m_i\rangle = \sqrt{m_i + 1}|m_i + 1\rangle$. Since we have three types of waves, it is convenient to abbreviate the tensor-product state $|m_1\rangle \otimes |m_2\rangle \otimes |m_3\rangle$ as $|m_1, m_2, m_3\rangle$. This number basis is natural for the quantized problem.

While the Schrödinger picture $i\partial_t|\psi\rangle = H|\psi\rangle$ is more convenient for quantum simulations, the connection between the quantum and classical models is more transparent in the Heisenberg picture $d_t a = i[H, a]$. For three-wave interactions, equations (2.1) are the Heisenberg equations from the cubic Hamiltonian

$$H_T = i g a_1^\dagger a_2 a_3 - i g^* a_1 a_2^\dagger a_3^\dagger. \tag{2.4}$$

A degenerate form of H_T , where $a_1 = a_2$, commonly arises in optomechanical systems (Aspelmeyer & Schwab 2008; Kong *et al.* 2018; Lake *et al.* 2020; Shang 2023). In our case, $a_1 \neq a_2$. The first term of H_T annihilates a seed and an idler photon to create a pump photon, while the second term of H_T is the reverse process where a pump photon decays into a seed and an idler photon. Although the Heisenberg equations for a_j are formally identical to the classical three-wave equations, the differences between the quantum and classical systems become apparent when one calculates higher-order cumulants. For example, because a_j and a_j^\dagger do not commute, the Heisenberg equation for n_i is different from its classical counterpart (Shi *et al.* 2021a). Similarly, for the four-wave interaction, (2.2) is the Heisenberg equation from the quartic Hamiltonian

$$H_F = -\frac{R}{2} a_2^\dagger a_2^\dagger a_2 a_2, \tag{2.5}$$

which is also known as the self-Kerr nonlinearity in the quantum literature (Kerr 1875). Since $R > 0$ for the modulational instability, the negative sign in H_F means that photons tend to condense together, which leads to a lower energy of the system. In (2.5), the operators are normal ordered $a_2^\dagger a_2^\dagger a_2 a_2 = n_2^2 - n_2$, which is different from other orderings such as $a_2^\dagger a_2 a_2^\dagger a_2 = n_2^2$. Different orderings differ by factors of the number operator n_2 .

For our purposes, the ordering does not matter, because the total Hamiltonian also includes the energy of the non-interacting harmonic oscillators $H_0 = \sum_j \omega_j a_j^\dagger a_j = \sum_j \omega_j n_j$. In the interaction picture of H_0 , these quadratic terms are trivially removed, and (2.4) and (2.5) should be understood as three- and four-wave interactions in the interaction picture of H_0 . Then, different orderings in (2.5) are equivalent up to a re-definition of ω_2 . Hence, it is sufficient to consider normal-ordered operators.

To use quantum Hamiltonian simulations to solve the quantized wave–wave interaction problems, we focus on the Schrödinger picture and use a basis that respects action conservation. In classical wave–wave interactions, $S_2 = n_1 + n_2$ and $S_3 = n_1 + n_3$ are known as the conserved wave actions. In the quantized model, $[H_T, S_2] = [H_T, S_3] = 0$, and H_F also commutes with S_2 and S_3 . Therefore, it is convenient to use eigenstates of S_2 and S_3 as the computational basis, which we call the action basis. For the laser pulse compression problem, since we are primarily interested in the seed wave a_2 , we label the action basis by

$$|\phi_j^{s_2, s_3}\rangle = |s_2 - j, j, s_3 - s_2 + j\rangle, \tag{2.6}$$

where j is the number of photons in the seed wave and is bounded within the range $j_{\min} \leq j \leq s_2$ with $j_{\min} = \max(0, s_2 - s_3)$. In other words, the $|\phi_j^{s_2, s_3}\rangle$ basis spans a $D = \min(s_2, s_3) + 1$ dimensional subspace of $|m_1, m_2, m_3\rangle$, where $m_1 = s_2 - j$, $m_2 = j$ and $m_3 = s_3 - s_2 + j$. The bounds for j come from the fact that $m_i \geq 0$. The non-negative integers s_2 and s_3 are eigenvalues of S_2 and S_3 , namely $S_2|\phi_j^{s_2, s_3}\rangle = s_2|\phi_j^{s_2, s_3}\rangle$ and $S_3|\phi_j^{s_2, s_3}\rangle = s_3|\phi_j^{s_2, s_3}\rangle$. In other words, the infinite-dimensional Hilbert space can be decomposed as a direct sum of finite-dimensional subspaces, where each subspace is labelled by a pair of quantum numbers (s_2, s_3) .

In the action basis, the Hamiltonian becomes block diagonal. The total Hamiltonian that governs the mixed three- and four-wave interaction problem is

$$H = H_T + H_F, \tag{2.7}$$

where $[H_T, H_F] \neq 0$. The matrix elements of H_T are $H_T|\phi_j^{s_2, s_3}\rangle = i g \eta_{j-1/2}^{s_2, s_3} |\phi_{j-1}^{s_2, s_3}\rangle - i g^* \eta_{j+1/2}^{s_2, s_3} |\phi_{j+1}^{s_2, s_3}\rangle$, where the reduced matrix element $\eta_{j-1/2}^{s_2, s_3} = \sqrt{(s_2 + 1 - j)j(s_3 - s_2 + j)}$ is inherited from the creation and annihilation operators. Notice that H_T only couples nearest neighbours in the action basis. Moreover, since the annihilation operator terminates at $m = 0$, the reduced matrix element vanishes for both bottom and top values of j . The bottom values of j are either $j = 0$ or $j = s_2 - s_3$, and in both cases $\eta_{-1/2}^{s_2, s_3} = \eta_{s_2 - s_3 - 1/2}^{s_2, s_3} = 0$. The top value is $j = s_2$, for which $\eta_{s_2 + 1/2}^{s_2, s_3} = 0$. In other words, in the $|\phi_j^{s_2, s_3}\rangle$ basis, the matrix of H_T is block tridiagonal, where each block is finite-dimensional, with no need for artificial truncation. Within each block, $H_F|\phi_j^{s_2, s_3}\rangle = -R \zeta_j |\phi_j^{s_2, s_3}\rangle$. The reduced matrix element is $\zeta_j = \frac{1}{2}j(j - 1)$, which is non-zero only when $j \geq 2$. This is intuitive because the self-Kerr effect is a photon–photon nonlinearity, so at least two photons are needed to see the effect. Because the matrix for H_F is diagonal, one can easily exponentiate the matrix analytically. Efficient quantum simulation algorithms also exist for diagonal Hamiltonian matrices with a smooth structure (Welch *et al.* 2014).

Because H is block diagonal in the action basis, we can perform Hamiltonian simulations separately in each (s_2, s_3) subspace. We span the wavefunction by

$$|\phi(t)\rangle = \sum_{s_2, s_3=0}^{\infty} \sqrt{p^{s_2, s_3}} \sum_{j=j_{\min}}^{s_2} c_j^{s_2, s_3}(t) |\phi_j^{s_2, s_3}\rangle. \tag{2.8}$$

The time-independent $p^{s_2, s_3} \geq 0$ is the probability that the state is within the (s_2, s_3) subspace, and the total probability $\sum_{s_2, s_3} p^{s_2, s_3} = 1$. Because different subspaces decouple,

the wavefunction is normalized separately within each subspace $\sum_j |c_j^{s_2, s_3}|^2 = 1$. After performing Hamiltonian simulations for the subspaces, which can be done in parallel, one can sum their contributions using weights that are predetermined by initial conditions.

For plasma problems, the initial states are often tensor-product states $|\phi(t=0)\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes |\psi_3\rangle$, where the three waves are initially unentangled. For each wave, $|\psi_i\rangle = \sum_m \beta_m^i |m\rangle$ can be expanded in its Fock basis, where $p_m^i = |\beta_m^i|^2$ is the probability that the i th wave occupies its m th state. For example, for classical laser pulse compression, the pump and seed lasers are initially in a coherent state $|\xi\rangle_c$, where $\xi = r e^{i\theta}$ is a complex number. For the coherent state, the probability amplitudes are $\beta_m = \exp(-\frac{1}{2}r^2)\xi^m/\sqrt{m!}$. In this case, p_m is a Poisson distribution, with $\langle n \rangle = r^2$ and $\langle n^2 \rangle = \langle n \rangle^2 + \langle n \rangle$, and p_m peaks near $m \sim r^2$. When $m \gg 1$, $p_m \simeq \exp(-r^2)(r^2/m)^m/\sqrt{2\pi m}$ decays faster than exponential. As another example, for interactions between quantum light, the initial state may be a squeezed vacuum state $|\xi\rangle_q$, for which $\beta_{2m+1} = 0$ for odd states and $\beta_{2m} = (\text{sech } r)^{1/2}(-\frac{1}{2} e^{i\theta} \tanh r)^m \sqrt{(2m)!}/m!$ for even states. In this case, $\langle n \rangle = \sinh^2 r$ and $\langle n^2 \rangle = 3\langle n \rangle^2 + 2\langle n \rangle$. The probability p_{2m} monotonically decreases, and $p_{2m} \simeq \text{sech } r (\tanh r)^{2m} / \sqrt{\pi m}$ decays much slower than a coherent state but faster than an exponential. In both examples, due to the superexponential decays of p_m , it is sufficient to keep track of a finite number of states in the Fock basis. Expanding the initial tensor-product state in the action basis gives

$$\sqrt{p^{s_2, s_3}} c_j^{s_2, s_3}(t=0) = \beta_{s_2-j}^1 \beta_j^2 \beta_{s_3-s_2+j}^3. \tag{2.9}$$

The simplest case is when only the pump wave is initially excited, which means that the seed and idler waves are initially in the ground state $\beta_m^2 = \beta_m^3 = \delta_{m,0}$. In this case, $p^{s_2, s_3} = |\beta_{s_2}^1|^2 \delta_{s_2, s_3}$ and $c_j^{s_2, s_3}(t=0) = \delta_{j,0}$. At later time, the wavefunction is spanned by $|\phi(t)\rangle = \sum_{s=0}^\infty \beta_s^1 \sum_{j=0}^s c_j^{s, s}(t) |\phi_j^{s, s}\rangle$. Another special case is when the idler is initially in the ground state, namely $\beta_m^3 = \delta_{m,0}$. In this case, $p^{s_2, s_3} = |\beta_{s_3}^1|^2 |\beta_{s_2-s_3}^2|^2 \Theta_{s_2, s_3}$, where the step function $\Theta_{i, j} = 0$ when $i < j$ and $\Theta_{i, j} = 1$ when $i \geq j$. The initial condition is $c_j^{s_2, s_3}(t=0) = \delta_{j, s_2-s_3}$, and at a later time the wavefunction is spanned by $|\phi(t)\rangle = \sum_{s_2=0}^\infty \sum_{s_3=0}^{s_2} \beta_{s_3}^1 \beta_{s_2-s_3}^2 \sum_{j=s_2-s_3}^{s_2} c_j^{s_2, s_3}(t) |\phi_j^{s_2, s_3}\rangle$. In both examples, the wavefunction initially occupies the ground state, namely the $j = j_{\min}$ state, within each (s_2, s_3) subspace.

For plasma simulations, the observables of interest are often the wave amplitudes, namely the expectation values $\langle n_i \rangle$ for the three waves. Due to action conservation, the three amplitudes are not independent. For given initial conditions, the conserved actions are $\langle S_2 \rangle = \sum_{s_2, s_3=0}^\infty p^{s_2, s_3} s_2$ and $\langle S_3 \rangle = \sum_{s_2, s_3=0}^\infty p^{s_2, s_3} s_3$. Because the probabilities p^{s_2, s_3} are time-independent, the expectation values $\langle S_2 \rangle$ and $\langle S_3 \rangle$ are constants of motion. The amplitudes of the pump and idler waves are related to the amplitude of the seed wave by

$$\langle n_1(t) \rangle = \langle S_2 \rangle - \langle n_2(t) \rangle, \tag{2.10a}$$

$$\langle n_3(t) \rangle = \langle S_3 \rangle - \langle S_2 \rangle + \langle n_2(t) \rangle. \tag{2.10b}$$

To compute the seed amplitude, we use the orthonormality of the action basis, which gives $\langle n_2 \rangle = \langle \phi | n_2 | \phi \rangle = \langle \phi | \sum_{s_2, s_3, j} \sqrt{p^{s_2, s_3}} c_j^{s_2, s_3} j | \phi_j^{s_2, s_3} \rangle = \sum_{s_2, s_3, j} p^{s_2, s_3} |c_j^{s_2, s_3}|^2 j$.

The amplitude can be computed in two steps:

$$\langle n_2^{s_2, s_3}(t) \rangle = \sum_{j=j_{\min}}^{s_2} j |c_j^{s_2, s_3}(t)|^2, \tag{2.11a}$$

$$\langle n_2(t) \rangle = \sum_{s_2, s_3=0}^{\infty} p^{s_2, s_3} \langle n_2^{s_2, s_3}(t) \rangle, \tag{2.11b}$$

where the first step computes the expectation value in each (s_2, s_3) subspace and the second step performs weighted sums using predetermined probabilities. For interactions involving quantum light, one may also be interested in observing higher-order cumulants, which can be measured in similar ways.

2.3. Exact dynamics

In the remaining part of this paper, we focus on a single subspace and drop the s_2 and s_3 superscripts to keep notations compact. When $s_2 = s_3 = 0$, the subspace is trivial because $D = 1$ is one-dimensional. For non-trivial dynamics, either s_2 or s_3 must be positive. Notice that although the underlying three- and four-wave interactions couple three and four photons, the dimension of the subspace $D = \min(s_2, s_3) + 1$ can take any integer value. The smallest non-trivial problem is $D = 2$, which requires a single qubit.

The exact dynamics involves two fundamental frequency scales g and R , from H_T and H_F , respectively (equations (2.4) and (2.5)). When $g = 0$, the dynamics is trivial because H_F is diagonal: under the influence of H_F alone, the occupation of $|\phi_j\rangle$ remains unchanged, and the dynamics is a pure phase precession. To change occupation numbers, a non-zero g is needed. Hence, for non-trivial dynamics, we can always normalize time by $\tau = |g|t$ and normalize the four-wave coupling by $\rho = R/|g|$. The Schrödinger equation becomes $i\partial_\tau \mathbf{c} = \mathbf{H}\mathbf{c}$, where the vector \mathbf{c} is the expansion coefficients such that $|\phi\rangle = \sum_{l=0}^{D-1} c_l |\phi_k\rangle$. The index $k = j_{\min} + l$, so c_0 is the probability amplitude that the seed wave is in the lowest occupied state and c_{D-1} is the probability amplitude that the seed wave is in the highest occupied state. The normalized matrix elements $H_{k',k} = (1/|g|)\langle \phi_{k'} | H | \phi_k \rangle$ are non-zero only along the diagonal and first off-diagonal bands:

$$H_{k,k} = -\frac{1}{2}\rho k(k-1), \tag{2.12a}$$

$$H_{k-1,k} = e^{i\theta} \sqrt{k(s_2 + 1 - k)(s_3 - s_2 + k)}, \tag{2.12b}$$

where $e^{i\theta} = ig/|g|$ and $H_{k,k-1} = H_{k-1,k}^*$. Notice that for the amplitude c_l , the index $k = j_{\min} + l$ may be shifted. For example, when $s_2 = 3$ and $s_3 = 4$, we have $D = 4$ and $j_{\min} = 0$, so $k = l = 0, 1, 2, 3$. On the other hand, when $s_2 = 4$ and $s_3 = 3$, even though we still have $D = 4$, now $j_{\min} = 1$, so $k = l + 1 = 1, 2, 3, 4$. In terms of l , we can rewrite $H_{k-1,k} = e^{i\theta} [l(D-l)(|s_3 - s_2| + l)]^{1/2}$, which is symmetric under $s_2 \leftrightarrow s_3$. The three-wave interaction is a phase-sensitive process because when $\varphi \neq \varphi'$, $|e^{i\varphi} + e^{i\varphi'}| < 2$. Here, the phase φ contains a contribution from θ , as well as a dynamical phase, which accumulates in time due to the diagonal four-wave interaction. If the relative phase between two adjacent levels changes, then the population transfer between them is reduced.

Since the Hamiltonian \mathbf{H} is time-independent, the exact dynamics is described by the unitary evolution operator $\mathbf{U}(\tau) = \exp(-i\mathbf{H}\tau)$, which can be obtained by direct diagonalization and exponentiation, at least for small problem sizes. For size $D = 4$, the exact behaviours of three examples are shown in figure 1. In all three examples, the coupling phase $\theta = 0$ and the constants of motion are $s_2 = 4$ and $s_3 = 3$, which means

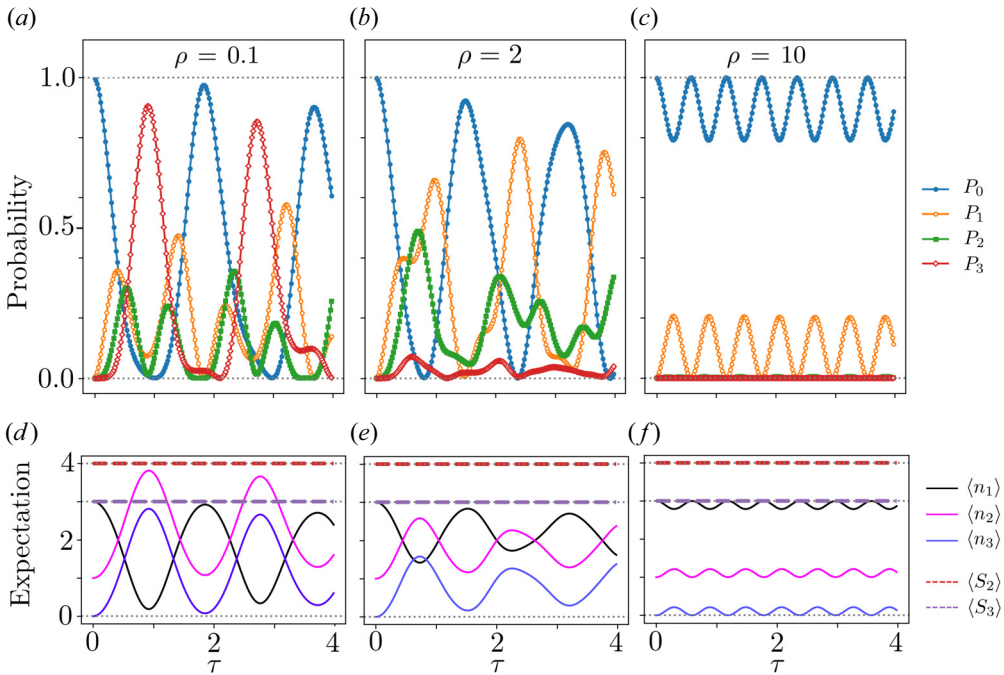


FIGURE 1. Exact dynamics of mixed three- and four-wave interaction problems in a $D = 4$ dimensional Hilbert space with constants of motion $s_2 = 4$ and $s_3 = 3$. Starting from the ground state, the probability amplitudes c are evolved in time, and the occupation probabilities $P_l = |c_l|^2$ ($a-c$) as well as the expected quanta in the three waves $\langle n_i \rangle$ ($d-f$) are computed on a classical computer. When $\rho = R/|g| = 0.1$, three-wave interaction dominates; when $\rho = 2$, three- and four-wave interactions compete; when $\rho = 10$, four-wave interaction dominates.

$j_{\min} = 1$ and $k = l + 1$, so $|\phi_k\rangle = |3 - l, 1 + l, l\rangle$. All examples start the time evolution from $c_0 = 1$ and $c_l = 0$ for $l = 1, 2, 3$ at $\tau = 0$. This ground state of the computational basis is by no means the ground state of the dynamical system. In fact, at $\tau = 0$, there are three photons in the pump wave a_1 and one photon in the seed wave a_2 , while the plasma wave a_3 is initially unexcited. At small τ , the three-wave interaction consumes a_1 to produce a_2 and a_3 , so the probabilities cascade from the ground state to higher states of the computational basis. At larger τ , the pump is depleted, so the inverse process dominates, where a_2 and a_3 merge into a_1 , and the probabilities cascade back to lower states. As can be seen from figure 1, the upward and downward cascades of probabilities repeat, but the dynamics is not periodic.

The dynamics is controlled by the dimensionless parameter ρ . When $\rho \ll 1$, as shown in figures 1(a) and 1(d), three-wave interaction dominates, which causes population transfer between the three waves. In this case, the much weaker four-wave interaction slowly accumulates phase mismatches that reduces the efficiency of population transfer, which is manifested by the decreasing oscillation amplitudes of $\langle n \rangle$ in figure 1(d). Here, the expected number of quanta in the three waves are calculated from the probability amplitudes c by $\langle n_1 \rangle = \sum_l (s_2 - j) |c_l|^2$, $\langle n_2 \rangle = \sum_l j |c_l|^2$ and $\langle n_3 \rangle = \sum_l (s_3 - s_2 + j) |c_l|^2$, where $j = j_{\min} + l$, and the summation is over $l = 0, \dots, D - 1$. From the above equations, it is clear that $\langle S_2 \rangle = s_2$ and $\langle S_3 \rangle = s_3$ are exact constants of motion, as marked by horizontal dashed lines in figure 1(d-f). In the opposite limit $\rho \gg 1$, as shown in figures 1(c) and 1(f), four-wave interaction dominates. In this case, the phases of

different $|\phi_j\rangle$ precess at drastically different rates, which inhibits population transfer because three-wave interaction requires phase matching. In this example ($\rho = 10$), the $0 \leftrightarrow 1$ transfer is strongly suppressed, while the $1 \leftrightarrow 2$ transfer, which accumulates phase mismatch at a greater rate, becomes nearly impossible. Finally, in the intermediate case $\rho \sim 1$, as shown in figures 1(b) and 1(e), three- and four-wave interactions compete, and the dynamics is more complicated. While four-wave interactions generate phase mismatches that suppress population transfer, three-wave interactions change the populations and affect how the phases are weighted. The intermediate cases are where simulations are most needed for predicting the behaviour of the system.

The expectation value $\langle n(\tau) \rangle$ can be measured efficiently using $O(\log D)$ circuits. Due to action conservation, the three expectation values $\langle n_1 \rangle = s_2 - \langle n_2 \rangle$, $\langle n_2 \rangle = j_{\min} + \langle \mathbf{O} \rangle$ and $\langle n_3 \rangle = s_3 - s_2 + \langle n_2 \rangle$ are derivable from

$$\langle \mathbf{O} \rangle = \sum_{l=0}^{2^n-1} |c_l|^2 = \frac{2^n - 1}{2} - \sum_{j=1}^n 2^{n-1-j} \langle \phi | \mathbf{Z}_j | \phi \rangle, \tag{2.13}$$

where we embed the D -level system into the Hilbert space of $n = \lceil \log_2(D) \rceil$ qubits. We identify $|l\rangle = |\phi_{j_{\min}+i}\rangle$ and map this state to $|q_1 \cdots q_n\rangle$ when $l = (q_1 \cdots q_n)_2$, where $q = 0$ or 1 and $(q_1 \cdots q_n)_2 = 2^{n-1}q_1 + \cdots + 2^0q_n$ is the binary representation of l . When $N = 2^n > D$, we pad the state vector with zeros, namely we set $c_l = 0$ for $l \geq D$. On the right-hand side of (2.13), $|\phi\rangle = \sum_l c_l |l\rangle$ is the quantum state, and \mathbf{Z}_j is the Pauli σ_z gate acting on the j th qubit. By measuring n single-qubit Pauli strings $\langle \mathbf{Z}_j \rangle$ for $j = 1, \dots, n$, (2.13) allows us to obtain $\langle \mathbf{O} \rangle$ using $n = O(\log D)$ measurement circuits with n steps of post-processing. To see why (2.13) holds, we start from $\mathbf{O} = \sum_l |l\rangle \langle l|$. In the computational basis, $\mathbf{O} = \text{diag}(0, 1, \dots, N - 1)$ is a diagonal matrix, which can be expanded in Pauli basis as $\mathbf{O} = \sum_l r_l \sigma_l^n$, where σ_l^n is a tensor product of n Pauli matrices. Explicitly, $\sigma_{(q_1 \cdots q_n)_2}^n := \bigotimes_{i=1}^n \sigma_{3q_i}$, where $\sigma_{3q} = \mathbf{I}$ when $q = 0$ and $\sigma_{3q} = \mathbf{Z}$ when $q = 1$. Using the fact that $\text{tr}(\sigma_i^n \sigma_j^n) = 2^n \delta_{ij}$, where δ_{ij} is the Kronecker delta, the expansion coefficient equals $r_l = 2^{-n} \text{tr}(\sigma_l^n \mathbf{O})$. To calculate the trace, we need diagonal elements of σ_l^n , and we denote its k th diagonal element by $[\sigma_l^n]_k$. By induction, one can show that $[\sigma_{(q_1 \cdots q_n)_2}^n]_{(b_1 \cdots b_n)_2} = \prod_{i=1}^n (-1)^{q_i b_i}$. Then, $\text{tr}(\sigma_{(q_1 \cdots q_n)_2}^n \mathbf{O}) = \sum_k [\sigma_{(q_1 \cdots q_n)_2}^n]_k [\mathbf{O}]_k = \sum_{b_1, \dots, b_n=0,1} [\sigma_{(q_1 \cdots q_n)_2}^n]_{(b_1 \cdots b_n)_2} (b_1 \cdots b_n)_2 = \sum_{b_1, \dots, b_n=0,1} \prod_{i=1}^n (-1)^{q_i b_i} \sum_{j=1}^n 2^{n-j} b_j = \sum_{j=1}^n 2^{n-j} \xi_j$. To calculate the sum $\xi_j = \sum_{b_1, \dots, b_n=0,1} \prod_{i=1}^n (-1)^{q_i b_i} b_j$, one can show by induction that if $q_i = 0 \forall i \neq j$, $\xi_j = 2^{n-1}(-1)^{q_j}$, otherwise $\xi_j = 0$. Consequently, only $n + 1$ traces are non-zero. The first non-zero trace is when $q_i = 0$ for all i , in which case $\sigma_{(0 \cdots 0)_2}^n = \mathbf{I}$ and $\text{tr}(\mathbf{O}) = 2^{n-1}(2^n - 1)$. Second, when $q_i = \delta_{ij}$ for a given j , $\sigma_{(0 \cdots 1 \cdots 0)_2}^n = \mathbf{Z}_j$ is a single-qubit gate and $\text{tr}(\mathbf{Z}_j \mathbf{O}) = -2^{n-1}2^{n-j}$. When more than one q is non-zero, the trace is zero. Therefore, we obtain $\mathbf{O} = 2^{-1}(2^n - 1)\mathbf{I} - \sum_{j=1}^n 2^{n-1-j} \mathbf{Z}_j$, from which (2.13) follows.

The exact quantum dynamics can be simulated efficiently using quantum Hamiltonian simulations (Berry *et al.* 2007, 2014; Low & Chuang 2017), exploiting the fact that our Hamiltonian matrix is 3-sparse. For example, the qubitization algorithm of Low & Chuang (2017) requires $O(1)$ ancilla qubits in addition to $O(\log D)$ qubits that encode the action basis. The query complexity of the qubitization algorithm, namely the number of terms one need to keep in the Jacobi–Anger expansion, is $O(\tau \|H\|_{\max} + \log(1/\epsilon) / \log \log(1/\epsilon))$, where ϵ is the desired precision and $\|H\|_{\max}$ is the matrix element with the largest absolute value. For plasma pulse compression, one is typically interested in attaining the maximum seed wave intensity. As shown in figure 1, the first peak of $\langle n_2 \rangle$ tends to be the highest. So, it is usually sufficient to simulate the dynamics for the first few oscillations. In other words, the maximum simulation time of interest is $\tau = O(1/|\lambda|_{\max})$,

where $|\lambda|_{\max}$ is the eigenvalue of H that has the largest absolute value. Therefore, the query complexity for the pulse compression problem is $O(\|H\|_{\max}/|\lambda|_{\max})$. In other problems, such as simulating the chaotic regime of wave-wave interactions, one might be interested in long-time dynamics, in which case the complexity may be higher. To estimate $\|H\|_{\max}$, using (2.12), the diagonal elements attain the maximum at the largest k value, which is $k = s_2$. When $s_2 \rightarrow \infty$, $\|H_{k,k}\|_{\max} \simeq \frac{1}{2}\rho s_2^2$. For off-diagonal elements, we express $H_{k,k-1}$ in terms of l . When $s_3 \rightarrow \infty$ at fixed s_2 , or $s_2 \rightarrow \infty$ at fixed s_3 , the maximum is attained at $l \simeq D/2$, where $\|H_{k,k-1}\|_{\max} \simeq (D/2)\sqrt{|s_2 - s_3| + D/2}$. When $s_2 = s_3 \rightarrow \infty$, the maximum is attained at $l \simeq 2D/3$ where $\|H_{k,k-1}\|_{\max} \simeq 2(D/3)^{3/2}$. Although $\|H\|_{\max}$ grows with D , the complexity remains $O(1)$ because $|\lambda|_{\max}$ also grows with D at a similar rate. To estimate the eigenvalues, consider two limits. (i) Consider the limit $\rho \rightarrow \infty$. When $s_2 > s_3$, because $j_{\min} > 0$, diagonal matrix elements dominate, and the eigenvalues are $\lambda_l \simeq H_{k,k} \rightarrow -\infty$, for $l = 0, \dots, D - 1$. When $s_2 \leq s_3$, because $j_{\min} = 0$, the first two diagonal elements are zero. The eigenvalues are $\lambda_0 \simeq |H_{0,1}|$, $\lambda_1 \simeq -|H_{0,1}|$ and $\lambda_l \simeq H_{k,k}$ for $l = 2, \dots, D - 1$, where $|H_{0,1}|^2 = (D - 1)(|s_2 - s_3| + 1) \ll \rho$. In both cases, $|\lambda|_{\max} \simeq \|H\|_{\max}$. (ii) Consider the limit $\rho \rightarrow 0$. Due to the absence of diagonal elements, the eigenvalues are either 0 or appear in $\pm\lambda_i$ pairs, where $i = 1, \dots, \lfloor D/2 \rfloor$. From the characteristic polynomial, $\sum_{i=1}^{\lfloor D/2 \rfloor} \lambda_i^2 = \sum_{l=1}^{D-1} |H_{k-1,k}|^2 = \frac{1}{12}D(D^2 - 1)(D + 2|s_2 - s_3|)$. Because $\sum_{i=1}^{\lfloor D/2 \rfloor} \lambda_i^2 \leq \lfloor D/2 \rfloor |\lambda|_{\max}^2$, we obtain a lower bound $|\lambda|_{\max} \gtrsim [\frac{1}{6}(D^2 - 1)(D + 2|s_2 - s_3|)]^{1/2}$. Comparing with expressions for $\|H_{k,k-1}\|_{\max}$, we see $\|H\|_{\max} = O(|\lambda|_{\max})$ when $D \rightarrow \infty$. For both limits, as well as for intermediate ρ , the complexity of quantum Hamiltonian simulation is $O(\|H\|_{\max}/|\lambda|_{\max}) = O(1)$, so the pulse compression problem can be simulated efficiently.

3. Implementing exact dynamics with error mitigation

The classical three- and four-wave interaction problems, when restricted to the temporal case with no spatial non-uniformity, are in fact not difficult to solve on classical computers. However, the quantum problem, which becomes important, for example, when the pump is in a squeezed state, becomes challenging for classical computers if the number of photons is large. The exact diagonalization of the 3-sparse $D \times D$ Hamiltonian takes $O(D)$ steps and $O(D^2)$ memory. For joule-class lasers typically used in plasma physics, the number of photons is of the order of 10^{19} , so exactly solving the quantum problem on classical computers is likely challenging. On the other hand, solving the quantum problem on future quantum computers will be efficient, which requires $O(\log D)$ qubits and $O(1)$ complexity. However, fault-tolerant quantum computers are not yet available. To push the limit of current devices and identify avenues for near-future improvements, we perform experiments on a superconducting device, using product-formula algorithms to approximate the exact dynamics.

We perform two-qubit experiments on Rigetti’s Aspen-M-3 processor (Caldwell *et al.* 2018), which is a superconducting device with multiple transmon qubits at a fixed topology with hardwired qubit–qubit couplings. The device is routinely calibrated to support single-qubit gates, as well as two-qubit gates like CZ and parametric XY(θ) gates. Each experiment is specified as a sequence of unitary operations, and each 4-by-4 unitary matrix is decomposed using Cartan decomposition into at most three two-qubit SQISW gates, sandwiched between single-qubit gates (Huang *et al.* 2023). The total gate sequence is executed on the hardware with the device initialized in the ground state. At the end of the gate sequence, the states of the two qubits are measured. The whole process of an experiment takes a few microseconds to run on hardware, with the overall time being dominated by a passive reset delay. We repeat each experiment for $M = 50\,000$ times




FIGURE 2. Occupation probabilities in a test problem with parameters $\rho = 2, \theta = 0, s_2 = 4$ and $s_3 = 3$. The blue lines are exact solutions from a classical computer, and the coloured symbols with error bars are results from the quantum device. When asked to enact the final unitary (black), the device performance is acceptable but not ideal. However, when asked to perform time evolution (orange), results on the device degrade to noise level after a few oscillations. The results are significantly improved using error-mitigation techniques (red), but the error bars grow exponentially.

to accumulate statistics for the final states, so that the shot noise, which scales as $O(1/\sqrt{M}) \sim 0.4\%$, is small compared with other sources of errors. At the end of M repeated experiments, we obtain a single data point along the time history of the evolution. Because projective measurement destroys quantum states, to obtain the next point along the time history, the simulation has to restart from the beginning in the form of a different experiment, which has its own sequence of unitary operations and is repeated another M times.

3.1. Implementation on a superconducting device

As the first test of the quantum device, we use it to enact the exact unitary operator. In figure 2, the solid blue lines are the exact occupation probabilities of the four states in our computational basis, which are computed using exact exponentiation on classical computers. The test problem uses parameters $\rho = 2, \theta = 0, s_2 = 4$ and $s_3 = 3$, which are identical to those of figure 1(b,e). The exact solutions serve as references for results on the quantum device. This first test is the simplest task that a quantum hardware can perform: for each time $\tau = N\Delta$, we compute the unitary exactly on a classical computer. The sequence of unitary operations for this experiment is thus constituted of just a single unitary, $U(N\Delta)$, and the results are shown in figure 2 as the black dashed lines. As can be seen from the figure, even when enacting a single dense unitary on the device, the fidelity is far from perfect. The important point is that because each time step uses the same gate depth, the performance does not degrade with τ , except perhaps at the very beginning where most population remains in the ground state and the unitary is near identity. In this test, because the gate sequence is so short, decoherence is not a leading cause of infidelity. Instead, most infidelity comes from coherent gate errors, in the sense that each gate realizes a slightly different unitary than what is intended.

In this simplest test, another source of error is readout, for which we have already corrected using an iterative Bayesian unfolding technique (Nachman *et al.* 2020). On the device level, the dispersive readout process is a scattering experiment, where a microwave pulse, whose frequency is off-resonant from the qubit transition frequency, is injected to interact with the qubit. The amplitude and phase of the returned microwave pulse

depend on the state of the qubit, and therefore enables a measurement of the qubit states. During the readout process, the state of the qubit may change, for example, due to intrinsic decoherence or induced interactions from the readout pulse. As a consequence, the true probability p_i that qubits are in state i may differ from the measured probability m_i . The vectors \mathbf{m} and \mathbf{p} are related by a response matrix \mathbf{R} by $\mathbf{m} = \mathbf{R}\mathbf{p}$, where \mathbf{R} is also called the confusion matrix, which is ideally close to the identity matrix. In practice, the matrix \mathbf{R} is constructed during calibrations, which prepare known quantum states using single-qubit gates, assuming they are ideal, and then immediately measure the qubit states. Due to statistical noise, directly inverting the matrix to find $\mathbf{p} = \mathbf{R}^{-1}\mathbf{m}$ often leads to artefacts, including $p \notin [0, 1]$ and $\sum_i p_i \neq 1$. These artefacts can be removed using iterative Bayesian unfolding $p_i^{n+1} = \sum_j (m_j R_{ji} p_j^n / \sum_k R_{jk} p_k^n)$, where n is the iteration number. When the iteration converges, the denominator $\sum_k R_{jk} p_k = m_j$ cancels with the numerator. Because the marginalized probability $\sum_j R_{ji} = 1$ for all i , $p_i^{n+1} = p_i^n$ converges to a stationary true value. In practice, starting from the initial guess $\mathbf{p}^0 = \mathbf{m}$, a few iterations are usually sufficient. Without readout error mitigation, results differ only slightly from the black dashed lines in figure 2, which suggests that the leading error is not readout error but rather coherent gate errors.

As the second test, we perform time evolution using the exact unitary $\mathbf{U}(\Delta)$, and the results are shown by the orange dashed lines in figure 2. In this set of experiments, $\mathbf{U}(\Delta)$ is compiled to native gates, and the gate sequence is repeated N times to enact $\mathbf{U}^N(\Delta)$. Because of the repetition, as $\tau = N\Delta$ increases, the gate depth increases linearly. The accumulation of errors leads to a degradation of fidelity, as can be seen from figure 2. The oscillation amplitudes decrease and $p(\tau)$ deviates further from the true solution as τ increases. At even larger τ values, the quantum states become fully scrambled, so $p \rightarrow 1/4$ approaches the fully mixed value for the four quantum states. Because $\mathbf{U}^N(\Delta)$ has a larger depth than $\mathbf{U}(N\Delta)$, the device performs worse in this test (orange lines) than in the previous test (black lines) as expected. The $\mathbf{U}^N(\Delta)$ results improve noticeably from Shi *et al.* (2021a) primarily because of the SQISW gate, which has significantly shorter duration and higher fidelity than the two-qubit gate used in our previous work.

3.2. Improving results with error suppression and mitigation

Because product formulas require even larger gate depth, we need to improve the results for the exact unitary before moving on to the next test. The dominant source of error on the quantum processor are two-qubit gate errors, which are typically an order of magnitude larger than single-qubit gate errors. On the Aspen architecture, this is not only due to a significantly longer two-qubit gate time, but also because activating the two-qubit gate requires tuning one of the qubits away from its optimal operating point so the qubit becomes more sensitive to flux noise (Didier 2019), which leads to a higher dephasing rate for the qubit.

We take a multi-pronged approach to minimizing two-qubit gate errors. First, the Aspen-M-3 chip used in our experiment has ~ 200 calibrated two-qubit gates available. We require only one of these for this circuit and are thus able to select high-performing candidates based on the reported fidelities. In our experiment, we perform this selection manually, but this optimization can be performed by compilers and is often referred to as addressing. Second, the Aspen chip offers both native CZ and XY(θ) gates, thus providing a choice of how to express our problem unitary. We observe that the XY family is particularly expressive (Peterson *et al.* 2020), allowing the expression of our target unitary using two XY($\pi/2$) gates and single-qubit gates (Huang *et al.* 2023). The parametric XY(θ) gate is composed of two pulses with a swap angle of $\pi/2$ while the variable angle

is achieved by a phase shift between the pulses (Abrams *et al.* 2020). However, we find that by using SQISW as our native gate, we can discard the second pulse from the XY gate, cutting the duration in half and reducing the two-qubit gate error by around 40%. Our native two-qubit gate is thus a 64 ns SQISW which is combined with single-qubit rotations to produce highly expressive native cycles. Our target unitary is expressible in either two or three of these clock cycles depending on parameters of the problem. Our combined approach to native gate selection, which takes into account both the fidelity and expressiveness of the native two-qubit gates, allows us to achieve sufficiently low error rates so that it is possible to reach relatively large depths.

While our calibration and gate selection techniques suppress errors, setting the fundamental performance ceiling of the circuit, a family of error-mitigation techniques allows us to tailor the noise, and recover unbiased estimates of observables. As mentioned earlier, a major source of error is coherent gate error. Coherent errors may be the result of either control errors, which should have ideally been calibrated away, or crosstalk, which are coherent errors caused by the state of nearby idling spectator qubits. Such errors can be suppressed using dynamical decoupling (Tripathi *et al.* 2022; Zhou *et al.* 2023; Evert *et al.* 2024). We do not attempt this here because our circuit does not have significant idle periods. Finally, we may suffer from unintended control errors from neighbouring control signals. All three classes of errors are major contributors to infidelity on current devices. Moreover, while long-term changes, such as temperature drifts that affect the control electronics, can be removed by routine calibrations, current devices also suffer from short-term and unpredictable changes. For example, chemical residues from fabrication process or material defects can interact with the qubits, which changes qubit frequencies and coherence times (Klimov *et al.* 2018; Müller, Cole & Lisenfeld 2019; Cho *et al.* 2023). As another example, when the device is struck by cosmic rays, whose energy is often higher than that of the superconducting gap, Cooper pairs are broken, creating quasiparticles, and thus changing the qubit frequency (Vepsäläinen *et al.* 2020). Due to these changes, a control pulse that is calibrated to enact a specific unitary for the original set of qubit parameters will now deviate from the intended unitary operation. The resultant coherent gate errors are particularly damaging because their behaviours are difficult to predict. Rather than adding up systematic errors in a simple way, quantum interference may cause the errors to transiently disappear, only to re-emerge at a later time. In comparison, errors due to decoherence, or more specifically due to depolarizing noise, behave in a much simpler and more predictable way: the depolarizing errors simply lead to an exponential decay of the Bloch vector, which can be easily modelled once the decay exponent is measured.

To mitigate coherent errors, we convert them into stochastic Pauli errors using a random compilation technique (Wallman & Emerson 2016; Hashim *et al.* 2020). The technique exploits the fact that the decomposition of a target unitary into elementary gates is not unique. Suppose one has found a particular decomposition $\mathbf{U} = (\prod_{k=1}^K \mathbf{C}_k \mathbf{G}_k) \mathbf{C}_0$, where \mathbf{C} are single-qubit gates and \mathbf{G} are two-qubit gates. Then, a different but equivalent decomposition can be constructed in two steps. First, the single-qubit gates are replaced by $\tilde{\mathbf{C}}_k = \mathbf{T}_{k+1} \mathbf{C}_k \mathbf{T}_k^c$ for $k = 0, \dots, K$, where \mathbf{T}_{k+1} is chosen at random from a subgroup of \mathbf{C} called the twirling group. An important property of the twirling group is that $\mathbf{T} \mathbf{G} = \mathbf{G} \mathbf{T}'$. In other words, a single-qubit gate in \mathbf{T} can be commuted across two-qubit gates to become another single-qubit gate. Due to this property, once \mathbf{T}_k is chosen, if one takes $\mathbf{T}_k^c = \mathbf{G}_k \mathbf{T}_k^\dagger \mathbf{G}_k^\dagger$ for $k = 1, \dots, K$, then $\mathbf{T}_k^c \mathbf{G}_k \mathbf{T}_k = \mathbf{G}_k$. Taking $\mathbf{T}_0^c = \mathbf{I}$ to be the identity, one thus finds another decomposition $\mathbf{U} = \mathbf{T}_{K+1}^\dagger (\prod_{k=1}^K \tilde{\mathbf{C}}_k \mathbf{G}_k) \mathbf{C}_0$. Second, to reduce unnecessary single-qubit gate depth, the new single-qubit operations $\tilde{\mathbf{C}}_k$, for $k = 1, \dots, K - 1$, and

$T_{K+1}^\dagger \tilde{C}_K$ are compressed and simplified to elementary single-qubit gates. In this paper, the two-qubit gate we use is the SQISW gate $G = \exp[-i(\pi/8)(X \otimes X + Y \otimes Y)]$, where X and Y are Pauli matrices. Notice that the twirling does not directly touch the two-qubit gates, which are hard to implement on hardware. The randomness of the compilation is introduced solely from the intermediate single-qubit gates, which are easy to adjust. Wallman & Emerson (2016) shows that if the Pauli twirling gates are chosen independently and if hardware Pauli errors are also independent, then random compilation transforms any gate errors into stochastic Pauli errors. In other words, as shown in the Appendix, suppose the errors of a quantum channel, when represented by the Pauli-transfer matrix, have off-diagonal components before twirling. Then, after twirling, the errors become purely diagonal, which means coherent interference of errors is removed. Because our native two-qubit gate is a non-Clifford gate, we cannot apply full Pauli twirling. Rather, we use a pseudo-twirling technique which tailors a smaller subset of coherent errors using the group of single-qubit rotations which can be successfully inverted by single-qubit gates. This twirling group is less powerful than Pauli twirling, but still tailors the noise effectively in most situations. The twirling is performed using the TrueQ software library (Beale *et al.* 2020). In our experiments, we construct the logical circuit and compute 50 random compilations. Each compilation has an identical pulse schedule, and thus an identical noise model. The randomization of single-qubit gates is performed by updating angles of our virtual Z gates. Such updates can be made with high efficiency, allowing a large number of randomizations of the circuit to be executed in quick succession.

The final step of error mitigation is to compensate for the suppression of observables using a rescaling technique (Ville *et al.* 2022). After twirling of a unitary operation U , the noise channel becomes approximately $\mathcal{E}(\rho) = \sum \lambda_P \mathbf{P} \rho \mathbf{P}^\dagger$, where the summation is over all tensor products of single-qubit Pauli operators \mathbf{P} . The coefficient λ_P , called the Pauli decay constant, is specific to the Pauli operator \mathbf{P} but is independent of the unitary U that is being performed. Because the λ_P values are bounded by their mean $\bar{\lambda}$ as $2\bar{\lambda} - 1 \leq \lambda_P \leq 1$, Ville *et al.* (2022) propose to use a single $\bar{\lambda}$ value to correct for all Pauli errors. This approximation becomes exact when the errors are fully depolarizing, which means that all Pauli channels decay in the same way. In this case, the measured expectation value \tilde{E} for any E of interest is given by a simple rescaling $\tilde{E} = \bar{\lambda}E$, because the length of the Bloch vector, which measures the purity of the state, shrinks by $\bar{\lambda}$. For example, in our two-qubit problem, the expected occupation of a state beyond its fully mixed value is $(\tilde{p} - \frac{1}{4}) = \bar{\lambda}(p - \frac{1}{4})$, where p is the occupation probability for a pure state after one unitary operation. Then, after N unitary operations, we can purify the probability by a rescaling $p = \frac{1}{4} + \frac{1}{\bar{\lambda}^N}(\tilde{p} - \frac{1}{4})$. In other words, after measuring the probability of a state, we subtract the noise $\frac{1}{4}$ and amplify the remaining signal exponentially by a rescaling factor $(1/\bar{\lambda})^N$. We provide justifications for using the simple rescaling technique in the Appendix. Notice that while amplifying the signal, this purification procedure also amplifies statistical error bars exponentially.

In practice, we estimate the mean value of Pauli decay constants $\bar{\lambda}$ using cycle error reconstruction (Erhard *et al.* 2019; Carignan-Dugas *et al.* 2023). Cycle error reconstruction is a technique for measuring the Pauli infidelities of a dressed cycle with multiplicative precision. A Pauli-twirled cycle results in a noise channel which is diagonal in the Pauli basis. After the qubits are prepared in an initial state, the cycle of interest is repeated M times and inverted to create an identity. By repeating the procedure at different values of M , we can extract a state-preparation and measurement (SPAM) robust error rate for the initial state. The ensemble of error rates allow us to reconstruct the Pauli error rates of the cycle and their orbital averages, as shown in the Appendix. Together,

this set of measurements provides not only an overall measurement of the fidelity, but structured information about the error channel which can be used in more advanced error-mitigation techniques. The technique we use is scalable with the number of qubits, provided that the noise is sparse, which is a reasonable assumption in superconducting qubit architectures. Empirically, the result converges with the square root of the sampling size. In our case, the gate of interest is the SQISW gate. For a system of two qubits, it is affordable to perform cycle error reconstruction over all possible Pauli channels. The SQISW is a non-Clifford gate, meaning that Pauli twirls cannot be inverted by a following pair of single-qubit rotations. However, for the purposes of characterization, this is not a problem. The inversions are propagated to the end of the benchmarking circuit, and then inverted with a final SQISW gate. This procedure of measuring the purification constant $\bar{\lambda}$ is performed each time we run a batch of experiments on the hardware.

After performing the above error-mitigation steps, the hardware results for our test problem are shown in [figure 2](#) by the dotted red lines. The mitigated results of the second test now closely track the exact solutions, and the test performs even better than the first test (black lines), which does not use any mitigation. While the mitigation significantly improves the signals, without noticeably increasing the hardware overhead, the price we pay is exponentially growing error bars. At even larger simulation depth, the error bars will become comparable to the signals, beyond which the simulations need to stop.

4. Testing product formulas with limited quantum resources

With sufficient simulation depth, we can now test the next level of quantum simulations, without assuming that the exact unitary is known. Predicting the exact dynamics on classical computers requires exact exponentiation of \mathbf{H} , which becomes expensive for large problem sizes. As a prototypical step of quantum simulation, we split $\mathbf{H} = \mathbf{H}_T + \rho\mathbf{H}_F$ into two non-commuting terms, where the normalized four-wave coupling ρ is pulled out from the definition of \mathbf{H}_F . On future quantum computers, if the total Hamiltonian cannot be simulated efficiently, but its components can, then product formulas may be employed to approximate the exact dynamics. In our case, \mathbf{H}_F is diagonal, so it is trivial to exponentiate on classical computers; \mathbf{H}_T is tridiagonal with zero diagonal elements, so its spectrum appears in $\pm\lambda$ pairs. Otherwise, \mathbf{H} is not intrinsically more difficult to exponentiate than \mathbf{H}_T on classical computers. We do not expect a future quantum computer to benefit from separating \mathbf{H} into \mathbf{H}_T and \mathbf{H}_F when performing quantum Hamiltonian simulations. Nevertheless, we perform the separation in order to test the performance of product formulas on current quantum devices.

4.1. Product formulas

For given problem parameters, we compile unitary matrices $\mathbf{U}_T(\tau) = \exp(-i\mathbf{H}_T\tau)$ and $\mathbf{U}_F(\tau) = \exp(-i\mathbf{H}_F\tau)$ to native gates using a Cartan decomposition (Smith *et al.* 2020). Then, we use \mathbf{U}_T and \mathbf{U}_F to approximate the exact \mathbf{U} . To first order, we approximate a single step with time step size $\tau = \Delta$ using the formula

$$\mathbf{U}_1(\Delta) = \mathbf{U}_T(\Delta)\mathbf{U}_F(\rho\Delta), \quad (4.1)$$

where the normalized four-wave coupling ρ appears as a scaling of time, which is convenient on quantum computing platforms that support parametric compiling. The error per step of the first order formula is $O(\Delta^2)$, whose prefactor is proportional to the norm of

the commutator $[\mathbf{H}_T, \mathbf{H}_F]$. To second order, we use the symmetric formula

$$\mathbf{U}_2(\Delta) = \mathbf{U}_T\left(\frac{\Delta}{2}\right) \mathbf{U}_F(\rho\Delta) \mathbf{U}_T\left(\frac{\Delta}{2}\right), \tag{4.2}$$

whose error per step is $O(\Delta^3)$. One could instead place \mathbf{U}_F on the outside. With either placement, when repeating \mathbf{U}_2 , the adjacent \mathbf{U} of the same kind can be merged to reduce the required number of operations. For example, using (4.2), we can simplify $\mathbf{U}_2^2(\Delta) = \mathbf{U}_T(\Delta/2) \mathbf{U}_F(\rho\Delta) \mathbf{U}_T(\Delta) \mathbf{U}_F(\rho\Delta) \mathbf{U}_T(\Delta/2)$. Implementing this unitary sequence thus requires compiling three different unitary matrices $\mathbf{U}_T(\Delta/2)$, $\mathbf{U}_T(\Delta)$ and $\mathbf{U}_F(\rho\Delta)$. In a similar spirit, higher-order product formulas require more types of unitary operations. For example, to third order, we use the product formula

$$\mathbf{U}_3(\Delta) = \mathbf{U}_T\left(\frac{7\Delta}{24}\right) \mathbf{U}_F\left(\frac{2\rho\Delta}{3}\right) \mathbf{U}_T\left(\frac{3\Delta}{4}\right) \mathbf{U}_F\left(-\frac{2\rho\Delta}{3}\right) \mathbf{U}_T\left(-\frac{\Delta}{24}\right) \mathbf{U}_F(\rho\Delta), \tag{4.3}$$

whose error per step is $O(\Delta^4)$. Notice that two of the six unitary operations above are evolving backward in time. This feature is common for high-order product formulas: in order to cancel higher-order errors, more operations are needed and the time step sizes become increasingly constrained such that negative values become necessary. To achieve even higher-order approximations, we use Suzuki’s symmetric recurrence formula (Suzuki 1990), which allows construction of \mathbf{U}_{k+2} from \mathbf{U}_k by

$$\mathbf{U}_{k+2}(\Delta) = \mathbf{U}_k^2(p\Delta) \mathbf{U}_k[(1 - 4p)\Delta] \mathbf{U}_k^2(p\Delta), \tag{4.4}$$

where $p = 1/(4 - 4^{1/(k+1)}) > 1/4$, so the middle step is an evolution backward in time. Suzuki’s formula has a self-similar structure once it is fully expanded in terms of elementary \mathbf{U}_T and \mathbf{U}_F operations. As the order k increases, \mathbf{U}_k becomes increasingly accurate with $O(\Delta^{k+1})$ errors per step, but the required number of elementary operations increases exponentially as $O(5^{k/2-1})$, which incurs a significant computational cost. Thus, in practice, high-order Suzuki formulas are rarely used and low-order formulas offer the best compromise of accuracy versus speed.

With error-mitigation techniques, we are able to run experiments on the quantum device for up to about 200 two-qubit gates. The gate depth is deep enough that we can begin to compare results of different product-formula algorithms, which are shown in figure 3. In this set of experiments, different product formulas are employed to reach the same final simulation time $\tau_f = 3$ at a fixed gate depth. We choose not to exhaust the maximum gate depth such that error bars at the final time remain small.

With a fixed gate budget, because lower-order algorithms (equations (4.1) and (4.2)) require fewer unitary operations per step, they can afford to use smaller time step sizes. In figure 3, both first-order (blue) and second-order (orange) algorithms yield results that match closely with the exact solutions (dashed lines), which are obtained on a classical computer by exact diagonalization and exponentiation of the total Hamiltonian matrix. Moreover, results on the quantum device (filled symbols) closely match results when the same product formula is used on a classical computer (open symbols).

In contrast, higher-order algorithms require significantly more unitary operations per step, and thus can only afford to use a much larger Δ for a fixed total gate depth. At third order (equation (4.3)), the time resolution is still sufficient to capture the oscillatory dynamics, and results on hardware are close to expected results (figure 3, green symbols). However, at fourth order (equation (4.4)), the coarse time step leads to a large discretization error. Although product-formula results (figure 3, red symbols) on

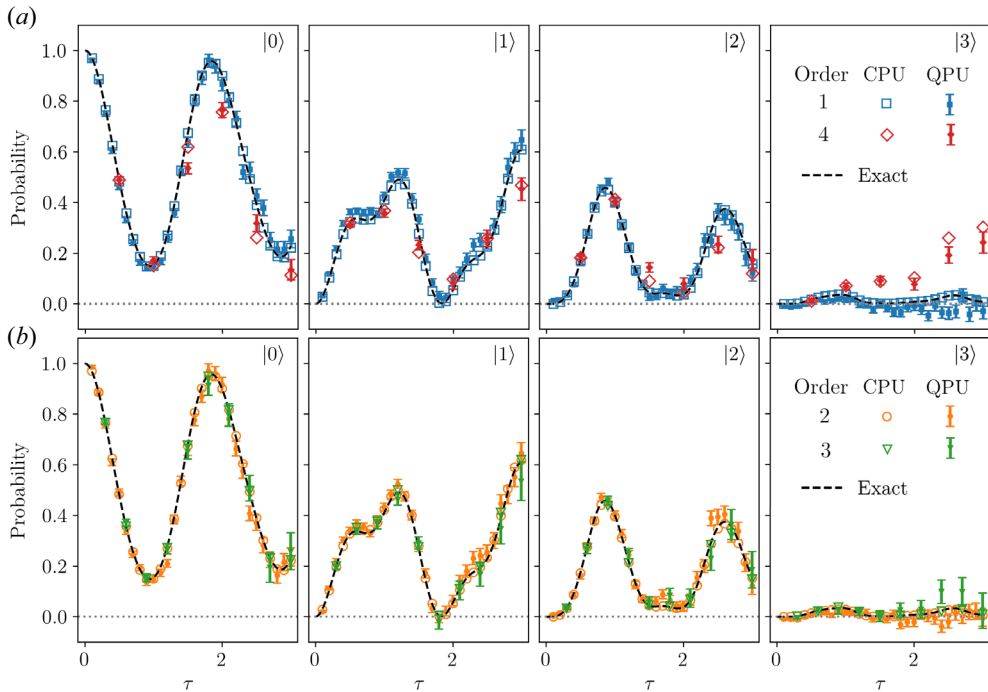


FIGURE 3. At a fixed gate depth, different product formulas are used for a test problem with parameters $\rho = 4$, $\theta = 0$, $s_2 = 3$ and $s_3 = 3$. For clarity, results from four cases are each split in two separate panels. The exact results are obtained by exponentiating the total Hamiltonian on a classical computer. Results of product formulas on a classical computer are shown by open symbols, and results obtained on quantum devices are shown by filled symbols with error bars. Because higher-order algorithms require more gates per step, for a fixed gate depth, they must use larger time step sizes to reach the same targeted final time $\tau_f = 3$, leading to larger discretization errors.

a quantum device remain close to results on a classical computer, the product formula no longer provides a good approximation to the exact dynamics, as can be seen from the deviations of the open red diamonds from the black dashed lines. The approximation becomes worse at even higher orders (not shown), which require exponentially more gates per step.

It is worth emphasizing that product formulas indeed become more accurate at higher orders, provided that the time step size Δ is fixed. In our tests, higher-order algorithms perform worse because Δ is changed, such that the total gate depth does not exceed what is viable on the quantum device. An analogy here is the run time on classical computers. While higher order algorithms are more accurate at a fixed resolution, they require more operations and therefore longer run time. When given a fixed run time, one is forced to use a coarser resolution, in which case higher order algorithms may perform worse than lower order algorithms.

4.2. Optimal use of limited quantum resources

On current quantum devices, which do not yet have operational error correction, the maximum gate depth is limited. To make the best use of the limited quantum resources, we can adjust the choice of algorithms and resolutions for a given problem. In our

case, the pulse compression problem seeks to determine the final seed laser intensity at the end of the interactions. The final time is set, for example, by the duration of laser pulses or the time to traverse the size of the mediating plasma. As a test problem, we fix $\tau_f = 1$ with parameter values $\rho = 4$, $\theta = 0$, $s_2 = 3$ and $s_3 = 3$. We perform Hamiltonian simulation on the quantum device using product formulas to evolve quantum states, and the measured occupation probabilities are post-processed to compute the expectation values of the three waves using (2.13). We measure the error of the simulation by $\epsilon = \{(1/N) \sum_{k=1}^N [n(k\Delta) - \langle n(k\Delta) \rangle]^2\}^{1/2}$, where n is the exact result on a classical computer and $\langle n \rangle$ is the expectation value obtained from the quantum hardware. In other words, we define the overall error of the simulation to be the 2-norm between the exact and measured time series, normalized by the number of time steps. The average error per step, ϵ , receives higher contributions from later steps of the simulation. In the definition of ϵ , it makes no difference whether we use n_1 , n_2 or n_3 , because $s_2 = n_1 + n_2$ and $s_3 = n_1 + n_3$ are exact constants of motion. Notice that the error of $\langle n \rangle$ is different from, albeit correlated with, the errors in the unitary \mathbf{U} . The unitary error, which is also known as the process infidelity, gives a more complete characterization of the hardware performance. But the expectation-value error is of more interest to the pulse compression problem: it is the same type of error that would typically be determined using a classical algorithm, and is much easier to measure than full process tomography in experiments.

The overall error receives contributions from two fundamental sources. First, algorithmic errors often arise when simulating a dynamical system with knowledge of only the exact solutions of its non-commuting subsystems. Algorithmic errors are unavoidable even on classical computers. In our test problem, we assume that the separate three- and four-wave unitary can be implemented exactly, and then use product formulas to approximate the total unitary. In this case, any finite time step size introduces a discretization error ϵ_Δ , which can be reduced either by using higher-order formulas at fixed Δ or by using the formula at a fixed order but with decreasing Δ . When using the Suzuki formula, demanding errors to scale as Δ^{q+1} per step requires M_q operations, which grows exponentially with q . On the other hand, to reach a target final time τ_f , the number of steps $N = \tau_f/\Delta$ increases only linearly when decreasing Δ . The total algorithmic error $\epsilon_1 = O(\tau_f \Delta^q)$ can in principle be made arbitrarily small by increasing q and decreasing Δ . However, in practice, given a limited run time, finite algorithm precision must be chosen. The trade-off between using a larger q versus a smaller Δ is strongly influenced by the second fundamental source of error: the hardware error. We measure hardware error by ϵ_Q , the error per unitary operation, which is analogous to round-off errors on classical computers. On future error-corrected quantum computers, it will be possible to suppress ϵ_Q to arbitrarily small values. However, on current noisy devices, with only error mitigation rather than error correction, ϵ_Q is substantial. Using randomized compilation, we transform coherent errors into random Pauli errors, which contribute to depolarizing noise together with intrinsic quantum decoherence. After error mitigation, the error for one operation becomes independent from that for the previous operation, so the total hardware error $\epsilon_2 = O(NM_q\epsilon_Q)$ accumulates linearly with the number of operations in the worst-case scenario. Because the algorithmic error ϵ_1 is independent of the hardware error ϵ_2 , the overall error is $\epsilon = \epsilon_1 + \epsilon_2$. Notice that ϵ_1 decreases with N , whereas ϵ_2 increases with N , so there is an optimal resolution Δ at which ϵ is minimized.

The trade-off between hardware and algorithmic errors is demonstrated by a suite of experiments, whose results are shown in figure 4. We test the four product formulas using a common test problem, whose parameters are $\rho = 4$, $\theta = 0$, $s_2 = 3$, $s_3 = 3$ and $\tau_f = 1$. For each order of the product formula, the overall error first decreases with N due to the




FIGURE 4. For a given problem, minimum error is obtained at the trade-off between algorithmic errors and hardware errors. All test problems use common parameters $\rho = 4, \theta = 0, s_2 = 3$ and $s_3 = 3$. The targeted final time $\tau_f = N\Delta = 1$ is fixed, so a finer resolution Δ requires more steps N , which means algorithmic errors decrease with N at the expense of accumulating more hardware errors. An optimal resolution exists, where the overall error is minimized. At higher order, the optimal N shifts towards lower resolution, and the minimum error does not improve with the algorithm order.

reduction of algorithmic errors at finer resolution. However, when N exceeds an optimal value N_* , the error starts to increase due to the accumulation of hardware errors. At small N , the resolution is too coarse to resolve the dynamics, so the errors are $O(1)$, which are comparable to the signals. If N_* had been larger, one would expect to see that ϵ for higher-order algorithms is smaller and decreases at a steeper slope. In our tests, because N_* is not large enough, such a behaviour is not clearly observed. At large N , where the accumulation of hardware errors dominates, ϵ increases roughly linearly with N for all orders. Because higher-order algorithms use more operations per step M_q , higher-order curves reside above lower-order ones in the log-log plots of ϵ , except between orders 1 and 2. Notice that although $M_1 = 2$ and $M_2 = 3$, after merging adjacent unitary operations of the same type, as discussed after (4.2), the first-order sequence has $2N$ operations, while the second-order sequence has $2N + 1$ operations, which is only slightly larger. The advantage of the second-order formula is not apparent because ϵ measures the error in n rather than in \mathbf{U} . On a classical computer, we observe a similar behaviour: at small N , while the second-order formula has a smaller error in \mathbf{U} , it has a larger error in n ; at large N , while the second-order formula has a larger error in \mathbf{U} , it has a smaller error in n . These behaviours are perhaps peculiar to our mixed three- and four-wave interaction problem. At intermediate N , higher-order algorithms tend to have a smaller optimal N_* . This behaviour can be understood from $\epsilon_1 \simeq C/N^q$ and $\epsilon_2 \simeq AB^qN$, where B^q comes from the exponential scaling of M_q . The minimum of $\epsilon(N)$ is reached at $N_* = (1/B)(Rq)^{1/(q+1)}$, where $R = BC/A$. The function $N_*(q)$ is not monotonic: it increases to a maximum before decreasing as $q^{1/q}$. When the hardware error is small compared with the algorithmic error, R is large, in which case the maximum of $N_*(q)$ is reached at a small q value. This means that N_* decreases with order, which is what we observe in our experiments.

While hardware errors are negligible in classical computers, current quantum devices have significant hardware errors. These errors make the trade-off with algorithmic errors a relevant issue. It is worth pointing out that when considering the effect of round-off errors, such a trade-off also exists for classical computers, but, due to the high precision available on today's classical computers, this is far less consequential. In the small

hardware error limit, $R = BC/A \rightarrow \infty$, so N_* monotonically decreases for $q \geq 1$. At a fixed q , N_* increases with R , which means a smaller hardware error can support the use of a finer resolution. On a typical classical computer, $A \sim 10^{-12}$ for our test problem, so that $N_* \sim 10^4$, at which the total error $\epsilon \sim 10^{-8}$ is negligible. In contrast, current quantum hardware has error $A \sim 10^{-2}$, which means in our test problem, the hardware can only support $N_* \sim 10^1$, for which the total error $\epsilon \sim 10^{-2}$ is noticeable. In such a scenario, there is no clear benefit of using higher-order product formulas because at a given hardware error, such formulas favour the use of lower resolution. But when the resolution becomes too low to resolve the dynamics, the overall error becomes worse with increasing algorithm order. In our test, higher-order algorithms do not perform better than a first-order algorithm at their respective optimal N_* . However, this may change on future quantum devices where ϵ_Q becomes even smaller. Similar conclusions have recently been reached for simulating the transverse-field Ising model and the XY model on noisy quantum computers (Avtandilyan & Pogosov 2024).

5. Conclusion

In summary, we develop a quantization approach for solving nonlinear wave–wave interaction problems on quantum computers. The approach becomes necessary when simulating interactions of quantum light with plasmas. The mixed three- and four-wave interaction problem serves as a non-trivial test for current quantum devices. We implement product-formula algorithms using two superconducting qubits along with a suite of error-mitigation techniques. The mitigated hardware error is small enough that we can begin to compare different product formulas and perform interesting simulations on quantum devices. In the future, when intrinsic hardware errors become even smaller, it may become feasible to explore more sophisticated algorithms for more non-trivial problems.

Acknowledgements

We thank B. Sundar and J. Andress for providing feedback to earlier versions of this paper.

Editor Nuno Loureiro thanks the referees for their advice in evaluating this article.

Funding

This work was performed under the auspices of the US Department of Energy Fusion Energy Sciences by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and under grant number SC-FES SCW1736. The paper is reviewed and released under LLNL-JRNL-865324.

Declaration of interests

The authors report no conflict of interest.

Appendix. Physically motivated error model

As discussed in § 3.2, we use twirling to convert coherent errors to depolarizing channels and rescale probabilities to compensate for exponential decays. In this Appendix, we provide justifications for these mitigation steps by investigating a physically motivated error model. We use the technique of cycle error reconstruction to measure an empirical error model of our gate under Pauli twirling. Although our native gate is non-Clifford, we are able to propagate the two-qubit inversion to the end of the circuit, meaning that the inversion appears only as a term in the SPAM coefficient of the decay. We perform the cycle error reconstruction for 1-body errors, meaning that we estimate the nine two-qubit

Pauli error rates (XX , XY , XZ , ...) and the six single-qubit Pauli error rates (XI , IX , ...). However, not all of these error rates are individually observable. The entangling gate induces orbitals of certain types of Pauli errors, meaning that we can only observe the product of the decay fidelities rather than individual Pauli errors.

The reconstructed Pauli error rates characterize the effective noise model under Pauli twirling with high accuracy. In order to determine a full Pauli error model, namely a model with a probability for each type of Pauli error, a method of distributing the degenerate error rates to their components is required. This can be achieved by simply assuming, for example, that both Pauli errors contribute to the product equally (van den Berg *et al.* 2022). Alternatively, one can perform further characterization to isolate the marginal effect of error channels (van den Berg & Wocjan 2024). The approach we take is to fit an open-system simulation to the observed Pauli error rates.

Using physical values of qubit parameters, including qubit frequencies, anharmonicities, a fixed coupling extracted from gate operation parameters and the parameters of our calibrated pulse, we are able to represent the full two-qubit system using a Lindblad model, as detailed in figure 5(a). In order to ensure the model replicates the system dynamics, we simulate our calibrated SQISW gate, which is enacted by a flux pulse on the tunable qubit. We adjust the qubit–qubit coupling to match our calibration. The determined coupling is 8.1 MHz, which is consistent with the design value and the operation of the two-qubit gates.

To simulate the effect of decoherence, we adjust damping (T_1) and dephasing (T_2) terms in the Lindblad model, starting with the reported values on the device. The simulation allows us to determine an effective superoperator for the SQISW gate. The errors, namely the distance from the superoperator of an ideal gate, are shown in the left-hand panel of figure 5(b). Using the simulation, we can also apply pseudo-twirling and Pauli twirling to determine the twirled superoperators, whose errors are shown in the middle and the right-hand panels of figure 5(b). The Pauli-twirled superoperator can then be transformed to Pauli error rates via a Walsh–Hadamard transform, which is compared with the observed error rates. Using reported coherence times, some important aspects of the observed Pauli error rates are reproduced, as shown in the left-hand panel of figure 5(c). However, the overall error rate is different and there are significant discrepancies in the error profile.

To improve our error model, we consider two additional effects. First, single-qubit gates also contribute noise to the cycle. Because single-qubit gates are randomized under twirling, we model this as a depolarizing error. Second, coherence decreases when qubits are under modulation (Didier 2019). Thus we allow the coherence times of the model to vary in order to match the observed error profile. The result is a closer match between the simulated and the observed error profiles, as shown in the right-hand panel of figure 5(c). The model converges on a significantly decreased T_1 time for both qubits. While T_1 is not normally expected to decrease under modulation, other mechanisms, such as leakage or two-level system loss, might have a similar effect during Pauli error reconstruction.

With a simulated error model in hand, we now have the ability to generate a Pauli error model that both reproduces the empirically observed Pauli error rates and allows us to decompose the errors into individual Pauli terms. While probabilistic error amplification (and probabilistic error cancellation strategies make use of these decompositions (van den Berg *et al.* 2022; Ferracin *et al.* 2022), our rescaling technique naively assumes depolarizing noise which uniformly decreases observable values. Moreover, in the characterization phase, we were able to use Pauli twirling by propagating the inversion of the benchmark circuit, but this is not possible in algorithm circuits. Thus, for our problem circuits, we use a pseudo-twirling technique which does not tailor noise as completely as Pauli twirling.

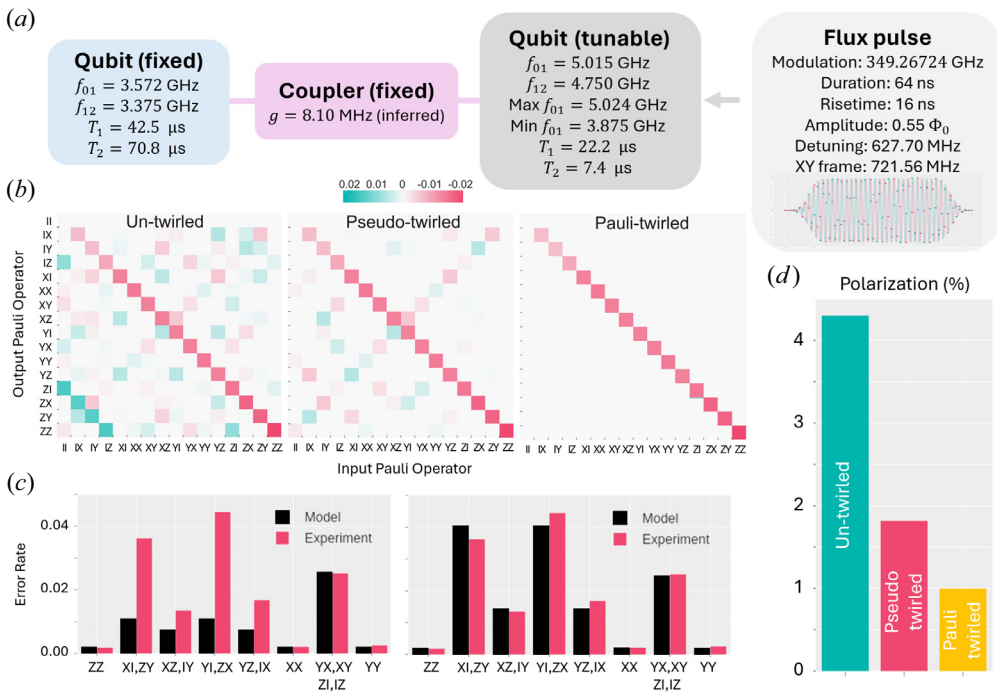


FIGURE 5. Modelling of our two-qubit SQISW gate reveals that errors become approximately depolarizing noise after twirling. (a) The two-qubit system we use consists of a fixed qubit and a tunable qubit, connected by a fixed coupler. We model the system using physical parameter values. The SQISW gate is enacted by a flux pulse that modulates the tunable qubit. (b) We compute errors in the Pauli transfer matrix for the SQISW gate. The error is entirely incoherent with 1.3% infidelity, but has a complex structure (left). After pseudo-twirling, the infidelity does not change, but the error is simplified and symmetrized (middle). After Pauli twirling, the same infidelity manifests only as errors along the diagonal (right). (c) After Pauli twirling, the observed Pauli errors for the SQISW gate (red) are not entirely explained by the model when using reported decoherence time T_1 and T_2 (left), which are measured when the gate is not in action. Nevertheless, by tuning the decoherence times, Pauli error reconstruction using our model (black) matches the observed errors (right). (d) We measure polarization of errors as the diamond-norm distance from pure depolarizing channels of the same infidelity. Twirling reduces noise polarization, which facilitates our error mitigation.

Using the physical noise model, we can now test the assumption that errors become depolarizing after twirling. As a metric of polarization, we compare the diamond-norm distance of the full superoperator, the Pauli-twirled superoperator and the pseudo-twirled superoperator with a depolarizing channel of the same infidelity. By this metric, a depolarizing channel would have a polarization of zero. As shown in figure 5(d), we find that the bare channel has a polarization of 4.3%, while the Pauli-twirled channel has a polarization of 1.0%. The pseudo-twirled channel has a polarization of 1.8%, meaning that, as expected, the pseudo-twirled channel is slightly less effective. The residual polarization implies that our rescaling technique has a systematic error, which depends on details of the problem, the observable and the noise. Nevertheless, we can conclude that for our physical noise model, Pauli twirling is effective in achieving a more depolarizing noise and that pseudo-twirling retains only moderately more polarization.

In summary, we develop an approach to fitting a physical noise model to a Pauli error reconstruction which can be efficiently measured. Fitting a physical noise model to these observable error rates provides a method of inferring a full noise model based on relatively few measurements. This model can be used for error-mitigation techniques, for error diagnostics and for studying the effectiveness of different twirling groups. By introducing even more effects into the simulation and constraining their values within reasonable bounds, we expect that better agreement with observations can be achieved, thus opening the door to further error-mitigation and error-suppression strategies. The approach outlined in this Appendix is scalable to cycles of any size, provided that the noise is relatively sparse. Both the Pauli noise reconstruction and simulation techniques become intractable for subsystem sizes larger than a few qubits, emphasizing the importance of limited crosstalk to the neighbourhood of a few qubits. Finally, we note that existing probabilistic error amplification and probabilistic error cancellation techniques typically rely on the circuit being executed with Pauli twirling, and work remains to generalize them to non-Clifford entangling gates, such as our SQISW gate.

REFERENCES

- ABRAMS, D.M., DIDIER, N., JOHNSON, B.R., DA SILVA, M.P. & RYAN, C.A. 2020 Implementation of XY entangling gates with a single calibrated pulse. *Nat. Electron.* **3** (12), 744–750.
- ASPELMEYER, M. & SCHWAB, K. 2008 Focus on mechanical systems at the quantum limit. *New J. Phys.* **10** (9), 095001.
- AVTANDILYAN, A.A. & POGOSOV, W.V. 2024 Optimal-order Trotter-Suzuki decomposition for quantum simulation on noisy quantum computers. [arXiv:2405.01131](https://arxiv.org/abs/2405.01131).
- BEALE, S.J., BOONE, K., CARIGNAN-DUGAS, A., CHYTROS, A., DAHLEN, D., DAWKINS, H., EMERSON, J., FERRACIN, S., FREY, V., HINCKS, I., *et al.* 2020 True-q. <https://doi.org/10.5281/zenodo.3945249>
- BERRY, D.W., AHOKAS, G., CLEVE, R. & SANDERS, B.C. 2007 Efficient quantum algorithms for simulating sparse Hamiltonians. *Commun. Math. Phys.* **270**, 359–371.
- BERRY, D.W., CHILDS, A.M., CLEVE, R., KOTHARI, R. & SOMMA, R.D. 2014 Exponential improvement in precision for simulating sparse Hamiltonians. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, pp. 283–292. ACM. <https://doi.org/10.1145/2591796.2591854>
- BOWEN, T.A., BADMAN, S., HELLINGER, P. & BALE, S.D. 2018 Density fluctuations in the solar wind driven by Alfvén wave parametric decay. *Astrophys. J. Lett.* **854** (2), L33.
- BOWEN, W.P. & MILBURN, G.J. 2015 *Quantum Optomechanics*. CRC Press.
- BREITENBACH, G., SCHILLER, S. & MLYNEK, J. 1997 Measurement of the quantum states of squeezed light. *Nature* **387** (6632), 471–475.
- BRILLOUIN, L. 1914 Light diffusion by a homogeneous transparent body. *C. R. Hebd. Seances Acad. Sci.* **158**, 1331–1334.
- CALDWELL, S.A., DIDIER, N., RYAN, C.A., SETE, E.A., HUDSON, A., KARALEKAS, P., MANENTI, R., DA SILVA, M.P., SINCLAIR, R., ACALA, E., *et al.* 2018 Parametrically activated entangling gates using transmon qubits. *Phys. Rev. Appl.* **10** (3), 034050.
- CARIGNAN-DUGAS, A., DAHLEN, D., HINCKS, I., OSPADOV, E., BEALE, S.J., FERRACIN, S., SKANES-NORMAN, J., EMERSON, J. & WALLMAN, J.J. 2023 The error reconstruction and compiled calibration of quantum computing cycles. [arXiv:2303.17714](https://arxiv.org/abs/2303.17714).
- CHO, Y., BECK, K.M., CASTELLI, A.R., WENDT, K.A., EVERT, B., REAGOR, M.J. & DUBOIS, J.L. 2023 Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits. [arXiv:2303.04261](https://arxiv.org/abs/2303.04261).
- DAVIDSON, R. 2012 *Methods in Nonlinear Plasma Theory*. Elsevier.
- DIDIER, N. 2019 Flux control of superconducting qubits at dynamical sweet spots. [arXiv:1912.09416](https://arxiv.org/abs/1912.09416).
- DODIN, I.Y. & STARTSEV, E.A. 2021 On applications of quantum computing to plasma simulations. *Phys. Plasmas* **28** (9).

- ENGEL, A., SMITH, G. & PARKER, S.E. 2021 Linear embedding of nonlinear dynamical systems and prospects for efficient quantum algorithms. *Phys. Plasmas* **28** (6).
- ERHARD, A., WALLMAN, J.J., POSTLER, L., METH, M., STRICKER, R., MARTINEZ, E.A., SCHINDLER, P., MONZ, T., EMERSON, J. & BLATT, R. 2019 Characterizing large-scale quantum computers via cycle benchmarking. *Nat. Commun.* **10** (1), 5347.
- EVERT, B., IZQUIERDO, Z.G., SUD, J., HU, H.-Y., GRABBE, S., RIEFFEL, E.G., REAGOR, M.J. & WANG, Z. 2024 Syncopated dynamical decoupling for suppressing crosstalk in quantum circuits. [arXiv:2403.07836](https://arxiv.org/abs/2403.07836).
- FERRACIN, S., HASHIM, A., VILLE, J.-L., NAIK, R., CARIGNAN-DUGAS, A., QASSIM, H., MORVAN, A., SANTIAGO, D.I., SIDDIQI, I. & WALLMAN, J.J. 2022 Efficiently improving the performance of noisy quantum computers. [arXiv:2201.10672](https://arxiv.org/abs/2201.10672).
- HANSEN, S.K., NIELSEN, S.K., SALEWSKI, M., STEJNER, M., STOBER, J., ASDEX UPGRADE TEAM, *et al.* 2017 Parametric decay instability near the upper hybrid resonance in magnetically confined fusion plasmas. *Plasma Phys. Control. Fusion* **59** (10), 105006.
- HASHIM, A., NAIK, R.K., MORVAN, A., VILLE, J.-L., MITCHELL, B., KREIKEBAUM, J.M., DAVIS, M., SMITH, E., IANCU, C., O'BRIEN, K.P., *et al.* 2020 Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor. [arXiv:2010.00215](https://arxiv.org/abs/2010.00215).
- HE, Y., QIN, H., SUN, Y., XIAO, J., ZHANG, R. & LIU, J. 2015 Hamiltonian time integrators for Vlasov-Maxwell equations. *Phys. Plasmas* **22** (12).
- HUANG, C., WANG, T., WU, F., DING, D., YE, Q., KONG, L., ZHANG, F., NI, X., SONG, Z., SHI, Y., *et al.* 2023 Quantum instruction set design for performance. *Phys. Rev. Lett.* **130** (7), 070601.
- JOSEPH, I. 2020 Koopman–von Neumann approach to quantum simulation of nonlinear classical dynamics. *Phys. Rev. Res.* **2** (4), 043102.
- JOSEPH, I., SHI, Y., PORTER, M.D., CASTELLI, A.R., GEYKO, V.I., GRAZIANI, F.R., LIBBY, S.B. & DUBOIS, J.L. 2023 Quantum computing for fusion energy science applications. *Phys. Plasmas* **30** (1).
- KERR, J. 1875 A new relation between electricity and light: dielectrified media birefringent. *Lond. Edinb. Dubl. Philos. Mag. J. Sci.* **50** (332), 337–348.
- KLIMOV, P.V., KELLY, J., CHEN, Z., NEELEY, M., MEGRANT, A., BURKETT, B., BARENDTS, R., ARYA, K., CHIARO, B., CHEN, Y., *et al.* 2018 Fluctuations of energy-relaxation times in superconducting qubits. *Phys. Rev. Lett.* **121**, 090502.
- KONG, C., LI, S., YOU, C., XIONG, H. & WU, Y. 2018 Two-color second-order sideband generation in an optomechanical system with a two-level system. *Sci. Rep.* **8** (1), 1060.
- KOUKOUTSIS, E., HIZANIDIS, K., VAHALA, G., SOE, M., VAHALA, L. & RAM, A.K. 2023 Quantum computing perspective for electromagnetic wave propagation in cold magnetized plasmas. *Phys. Plasmas* **30** (12).
- LAKE, D.P., MITCHELL, M., SANDERS, B.C. & BARCLAY, P.E. 2020 Two-colour interferometry and switching through optomechanical dark mode excitation. *Nat. Commun.* **11** (1), 2208.
- LIN, Y.T., LOWRIE, R.B., ASLANGIL, D., SUBAŞI, Y. & SORNBORGER, A.T. 2022 Koopman von Neumann mechanics and the Koopman representation: a perspective on solving nonlinear dynamical systems with quantum computers. [arXiv:2202.02188](https://arxiv.org/abs/2202.02188).
- LIU, J.-P., KOLDEN, H.Ø., KROVI, H.K., LOUREIRO, N.F., TRIVISA, K. & CHILDS, A.M. 2021 Efficient quantum algorithm for dissipative nonlinear differential equations. *Proc. Natl Acad. Sci. USA* **118** (35), e2026805118.
- LOUDON, R. 2000 *The Quantum Theory of Light*. OUP.
- LOW, G.H. & CHUANG, I.L. 2017 Optimal Hamiltonian simulation by quantum signal processing. *Phys. Rev. Lett.* **118** (1), 010501.
- MALKIN, V.M., SHVETS, G. & FISCH, N.J. 1999 Fast compression of laser beams to highly overcritical powers. *Phys. Rev. Lett.* **82** (22), 4448.
- MAY, M.Q. & QIN, H. 2023 Quantum three-wave instability. *Phys. Rev. A* **107**, 062204.
- MICHEL, P. 2023 *Introduction to Laser-Plasma Interactions*. Springer Nature.
- MORRISON, P.J. 2017 Structure and structure-preserving algorithms for plasma physics. *Phys. Plasmas* **24** (5).

- MÜLLER, C., COLE, J.H. & LISENFELD, J. 2019 Towards understanding two-level-systems in amorphous solids: insights from quantum circuits. *Rep. Prog. Phys.* **82** (12), 124501.
- NACHMAN, B., URBANEK, M., DE JONG, W.A. & BAUER, C.W. 2020 Unfolding quantum computer readout noise. *NPJ Quantum Inf.* **6** (1), 84.
- NAZARENKO, S. 2011 *Wave Turbulence*. Springer.
- PETERSON, E.C., CROOKS, G.E. & SMITH, R.S. 2020 Fixed-depth two-qubit circuits and the monodromy polytope. *Quantum* **4**, 247.
- RAMAN, C.V. & KRISHNAN, K.S. 1928 The negative absorption of radiation. *Nature* **122** (3062), 12–13.
- SHANG, C. 2023 Coupling enhancement and symmetrization of single-photon optomechanics in open quantum systems. [arXiv:2302.04897](https://arxiv.org/abs/2302.04897).
- SHI, Y., CASTELLI, A.R., WU, X., JOSEPH, I., GEYKO, V., GRAZIANI, F.R., LIBBY, S.B., PARKER, J.B., ROSEN, Y.J., MARTINEZ, L.A., *et al.* 2021a Simulating non-native cubic interactions on noisy quantum machines. *Phys. Rev. A* **103** (6), 062608.
- SHI, Y., QIN, H. & FISCH, N.J. 2017 Three-wave scattering in magnetized plasmas: from cold fluid to quantized Lagrangian. *Phys. Rev. E* **96** (2), 023204.
- SHI, Y., QIN, H. & FISCH, N.J. 2021b Plasma physics in strong-field regimes: theories and simulations. *Phys. Plasmas* **28** (4), 042104.
- SMITH, R.S., PETERSON, E.C., SKILBECK, M.G. & DAVIS, E.J. 2020 An open-source, industrial-strength optimizing compiler for quantum programs. *Quant. Sci. Technol.* **5** (4), 044001.
- SUZUKI, M. 1976 Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Commun. Math. Phys.* **51** (2), 183–190.
- SUZUKI, M. 1990 Fractal decomposition of exponential operators with applications to many-body theories and monte carlo simulations. *Phys. Lett. A* **146** (6), 319–323.
- TRIPATHI, V., CHEN, H., KHEZRI, M., YIP, K.-W., LEVENSON-FALK, E.M. & LIDAR, D.A. 2022 Suppression of crosstalk in superconducting qubits using dynamical decoupling. *Phys. Rev. Appl.* **18** (2), 024068.
- TROTTER, H.F. 1959 On the product of semi-groups of operators. *Proc. Am. Math. Soc.* **10** (4), 545–551.
- VAN DEN BERG, E., MINEV, Z.K., KANDALA, A. & TEMME, K. 2022 Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors. [arXiv:2201.09866](https://arxiv.org/abs/2201.09866).
- VAN DEN BERG, E. & WOCJAN, P. 2024 Techniques for learning sparse Pauli-Lindblad noise models. [arXiv:2311.15408](https://arxiv.org/abs/2311.15408).
- VEPSÄLÄINEN, A.P., KARAMLOU, A.H., ORRELL, J.L., DOGRA, A.S., LOER, B., VASCONCELOS, F., KIM, D.K., MELVILLE, A.J., NIEDZIELSKI, B.M., YODER, J.L., *et al.* 2020 Impact of ionizing radiation on superconducting qubit coherence. *Nature* **584** (7822), 551–556.
- VILLE, J.-L., MORVAN, A., HASHIM, A., NAIK, R.K., LU, M., MITCHELL, B., KREIKEBAUM, J.-M., O'BRIEN, K.P., WALLMAN, J.J., HINCKS, I., *et al.* 2022 Leveraging randomized compiling for the quantum imaginary-time-evolution algorithm. *Phys. Rev. Res.* **4**, 033140.
- WALLMAN, J.J. & EMERSON, J. 2016 Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A* **94**, 052325.
- WELCH, J., GREENBAUM, D., MOSTAME, S. & ASPURU-GUZIK, A. 2014 Efficient quantum circuits for diagonal unitaries without ancillas. *New J. Phys.* **16** (3), 033040.
- ZHAKAROV, V.E., L'VOV, V.S. & FALKOVICH, G. 1992 *Kolmogorov Spectra of Turbulence I, Wave Turbulence*. Springer.
- ZHOU, Z., SITLER, R., ODA, Y., SCHULTZ, K. & QUIROZ, G. 2023 Quantum crosstalk robust quantum control. *Phys. Rev. Lett.* **131** (21), 210802.