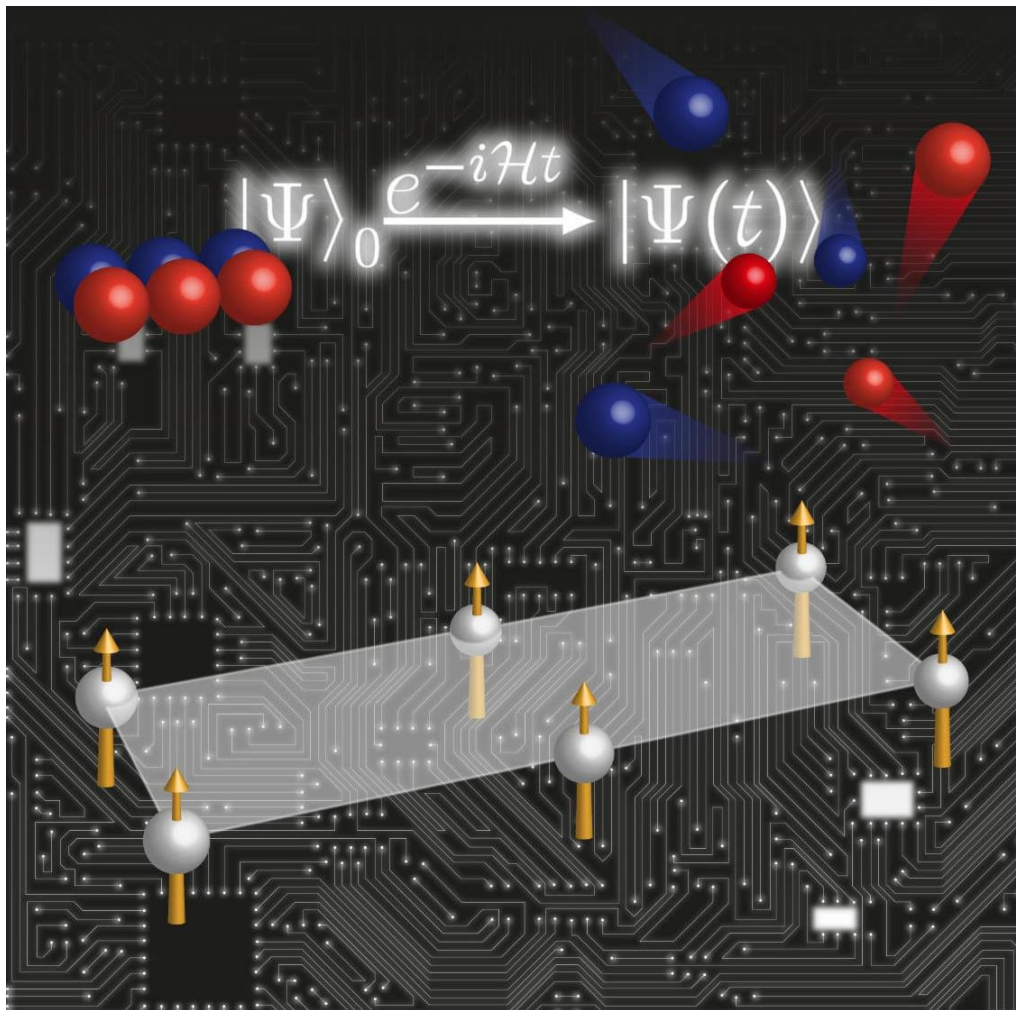


Digital quantum simulations and machine learning on near-term quantum processors

Francesco Tacchino



Tesi per il conseguimento del titolo



Università degli Studi di Pavia
Dipartimento di Fisica

DOTTORATO DI RICERCA IN FISICA – XXXII CICLO

Digital quantum simulations and machine
learning on near-term quantum processors

Francesco Tacchino

Submitted to the Graduate School of Physics in partial
fulfillment of the requirements for the degree of

DOTTORE DI RICERCA IN FISICA

DOCTOR OF PHILOSOPHY IN PHYSICS

at the

University of Pavia

Supervisor: Prof. Dario Gerace

Cover: Artistic representation of a digital quantum simulation.

**Digital quantum simulations and machine learning on near-term
quantum processors**

Francesco Tacchino

PhD thesis - University of Pavia

Pavia, Italy, September 2019

*A mamma e papà,
per il lavoro e l'esempio*

Abstract

Quantum mechanics is the gateway towards novel and potentially disruptive approaches to scientific and technical computing. In this thesis we explore, from a number of different perspectives, the effects of such strong relationship between the physical nature of information and the informational side of physical processes, with a focus on the digital quantum computing paradigm.

After an extensive introduction to the theory of universal quantum simulation techniques, we review the main achievements in the field and, in parallel, we outline the state of the art of near-term architectures for quantum information processing. We then move on to present novel and scalable procedures for the study of paradigmatic spin models on intermediate-scale noisy quantum processors. Through an innovative combination of quantum algorithms with classical post-processing and error mitigation protocols, we demonstrate in practice the full digital quantum simulation of spin-spin dynamical correlation functions, reporting experimental results obtained on superconducting cloud-accessible IBM Q devices. We also exhibit a practical use-case by successfully reproducing, from quantum computed data, cross section calculations for four-dimensional inelastic neutron scattering, a common tool employed in the analysis of molecular magnetic clusters.

The central part of the thesis is dedicated to the exploration of perspective hardware solutions for quantum computing. As it is not yet clear whether the currently dominant platforms, namely trapped ions and superconducting circuits, will eventually allow to reach the final goal of a fully-fledged architecture for general-purpose quantum information processing, the search for alternative technologies is at least as urgent as the improvement of existing ones or the development of new algorithms. After providing an overview of some recent proposals, including hybrid set-ups, we introduce quantum electromechanics as a promising candidate platform for future realizations of digital quantum simulators and we predict competitive performances for an elementary building block featuring nanomechanical qubits integrated within superconducting circuitry.

In the final part, we extend the reach of quantum information protocols beyond its traditional areas of application, and we account for the birth and rapid development of Quantum Machine Learning, a discipline aimed at establishing a productive interplay between the parallel revolutions brought about by quantum computing and artificial intelligence. In particular, we describe an original procedure for implementing, on a quantum architecture, the behavior of binary-valued artificial neurons. Formally exact and platform-independent, our data encoding and processing scheme guarantees in principle an exponential memory advantage over classical counterparts and is particularly well suited for pattern and image recognition purposes. We test our algorithm on IBM Q quantum processors, discussing possible training schemes for single nodes and reporting a proof-of-principle demonstration of a 2-layer, 3-neuron feed-forward neural network computation run on 7 active qubits. The latter is, in terms of the total size of the quantum register, one of the largest quantum neural network computation reported to date on real quantum hardware.

Introduction	vii
1 Digital quantum simulations	1
1.1 Theory of quantum simulations	3
1.1.1 From Feynman to Lloyd	4
1.1.2 Quantum simulations cookbook	5
1.1.3 Spin mapping and circuit decompositions	7
1.1.4 Suzuki-Trotter approximation and digital error	13
1.1.5 Extracting physical observables	16
1.1.6 Examples	18
1.2 State-of-the-art in digital quantum simulations	21
1.2.1 Trapped Ions	22
1.2.2 Superconducting circuits	25
1.2.3 Summary and comparison	27
2 Quantum dynamics on near-term quantum processors	31
2.1 The IBM Quantum Experience	31
2.2 Digital quantum simulations of spin Hamiltonians	36
2.2.1 Tunneling of $S = 1$ total magnetization	36
2.2.2 Spin-1/2 chains	38
2.2.3 Transverse Field Ising model	40
2.3 Dynamical quantum correlations	41
2.3.1 Quantum algorithms for dynamical correlation functions	41
2.3.2 Results and error mitigation	44
2.3.3 Error analysis and scalability	48
2.3.4 Four-dimensional inelastic neutron scattering	53
2.3.5 Entanglement and concurrence	57
3 Electromechanical Quantum Simulators	61
3.1 Perspective platforms for quantum information processing	61
3.2 Hybrid quantum technologies	63
3.2.1 Mechanical quantum devices	63
3.2.2 Quantum electromechanics	65
3.3 Electromechanical qubits for digital quantum simulations	67
3.3.1 Hybrid electromechanical building block	67
3.3.2 Effective qubit-qubit interactions	70
3.3.3 Single and two-qubit gates	71
3.3.4 Non-linearity of electromechanical qubits	76
3.3.5 Digital quantum simulations	82

4	Quantum artificial neural networks	85
4.1	The rise of quantum machine learning	85
4.1.1	Quantum subroutines for linear algebra	85
4.1.2	Quantum classification	87
4.1.3	Searching for quantum neurons	89
4.1.4	Perspectives on Quantum Machine Learning	92
4.2	A quantum artificial neuron	93
4.2.1	Quantum algorithm for binary-valued artificial neurons	94
4.2.2	Proof-of-principle realization on NISQ processors	101
4.2.3	Hybrid training scheme	103
4.3	A quantum feed-forward neural network	105
4.3.1	Connecting quantum artificial neurons	105
4.3.2	Pattern classification on a real quantum processor	107
	Conclusions	113
	A Additional data and results	115
A.1	Suzuki-Trotter decomposition	115
A.2	Digital quantum simulations of dynamical correlations	116
A.3	Fitting dynamical correlations	125
	List of Publications	147
	Acknowledgments	149

Introduction

Over slightly more than a century, quantum mechanics revolutionized the physical sciences like almost no other theory ever did. The advent and fortunes of modern optics, chemistry, solid state and particle physics, just to name a few examples, all relate to the deep principles governing the microscopic world. At the same time, foundational and interpretational problems connected to quantum theory challenged, and still challenge, the very meaning that we assign to scientific experiences: Bell's inequalities [1], together with a number of similarly counterintuitive results, certified the untenability of the orthodox combination of realism and locality, thus paving the way for an information-oriented vision of the universe [2]. Driven by that awareness, unveiling the physical nature of information [3, 4] opened new philosophical horizons: indeed, if on one hand computation and communication should ultimately be regarded as physical processes subject to an appropriate set of laws and limitations, it might also be productive to symmetrically frame all physical entities in operational terms, with states and events being placeholders for the reality of protocols, preparations and observations. In this spirit, we may then acknowledge that the way we speak about the world is at least as important as what we talk about: to say it in different words, information manipulation and the understanding of natural phenomena are unequivocally and intimately related.

The interplay between physical systems and their possible role as either computers or subjects of a computation is particularly strong, for conceptual and technological reasons, in the quantum realm: in this work, we will examine such relationship from different angles. The fundamentals of quantum computer science have been elaborated in the past few decades [5], proving mathematically the advantages offered by quantum information processing over classical counterparts. Within the same period of time, only preliminary experimental realizations of quantum computing architectures, and no practically relevant speed-up, were reported: genuine quantum mechanical properties, specifically coherence and entanglement, are fragile and elusive almost as much as they are potentially powerful. However, and despite the remarkable distance that, without any doubt, still separates us from the fully quantum-powered era of computation, a new wave of optimism is rising in the scientific community. Quantum technologies, from sensing to cryptography, are receiving increasing attention from institutional and industrial pioneers, and quantum computing remains at the forefront of those efforts. As classical computers both created new areas of research and enabled some of the most spectacular achievements of mankind, not least the landing on the Moon of which we celebrate in 2019 the 50th anniversary, the hope for a similar or even greater impact looks reasonable for the field of quantum information. Apart from a few much sought milestones, and some inevitable inflated claims, an incremental approach towards the final goal of fault-tolerant quantum computation seems nowadays the most viable strategy: in this quest, the so called Noisy Intermediate-Scale Quantum (NISQ) processors are certainly in the best position to play a vital role [6]. Indeed, finding early applications to practical problems, while keeping a steady progress on fundamental issues, is considered the key to a smooth and sustainable growth.

As we will see in Chapter 1, a very natural test bed for quantum computing is represented by quantum simulations. The importance of the latter, recalling the beginning of our discussion, can hardly be overstated: being able to simulate, namely describe and control, physical events represents one of the sharpest forms of scientific investigation. Moreover, it is well known that, when the

target model is quantum mechanical, the classical resources required to carry out such a task scale exponentially with the size of the problem, thus calling for a radical change in the numerical techniques to be adopted. We will concentrate upon the *digital* approach to quantum simulations, where the name itself manifests the focus on an abstract, discrete and platform-independent encoding and manipulation of information: within this scheme, any architecture designed for general-purpose, qubit-based quantum computing can equally well be regarded as a Universal Quantum Simulator capable of reproducing efficiently, under a suitable mapping, the dynamics of a large class of physical systems. After a formal introduction to digital quantum simulations and an extensive discussion about the most common algorithmic procedures and tools, we will then review the major advancements in the field and, in doing so, we will essentially assess the present status of quantum computing technologies. Special attention will be devoted to trapped ions and superconducting platforms, which are currently leading the way.

In Chapter 2, we will deploy digital quantum simulation protocols on real NISQ superconducting processors and we will present original and potentially scalable procedures to extract physically meaningful results from such non error-corrected quantum chips. After some preliminary tests on archetypal spin models, we will concentrate on the quantum simulation of dynamical spin-spin correlation functions, from which several energetic and structural characteristics of a given system can be extracted. The practical experiments that we are going to discuss, involving quantum registers containing up to 5 active qubits, were all performed by accessing remotely the real quantum backends made available by IBM via the IBM Quantum Experience program, using the Qiskit (Quantum Information Science Kit) open source python library. Besides providing explicit examples of the natural capabilities of such state-of-the-art quantum processors, we will conduct a thorough analysis of the actual noise sources and systematic errors affecting their performances. In particular, we will argue that, for digital quantum simulation purposes, systematic coherent inaccuracies are currently the limiting factor, and we will introduce and successfully employ an *a posteriori* error mitigation procedure based on general symmetry considerations and sum rules. A practical application of the results to a real physical problem, namely the calculation of the cross section for four-dimensional inelastic neutron scattering investigations of magnetic clusters, will also be described.

Having reviewed and demonstrated the current level of maturity of the dominant quantum technological platforms, in Chapter 3 we will move on to discuss perspective and alternative hardware solutions for quantum information processing. Focusing on the so called *hybrid* paradigm, in which different components with complementary strengths and weaknesses are combined to optimize the desired functionality, we are going to present an original proposal of an electromechanical architecture for digital quantum simulations. We will describe an intrinsically scalable elementary building block featuring, as quantum information carriers, two nanomechanical oscillators integrated within usual superconducting circuitry. Through numerical simulations, we will investigate the potential fidelities of a universal set of single- and two-qubit gates, which greatly benefit from the design of the set-up and from the promising coherence properties of nanomechanical devices. Moreover, we will extensively comment on a novel method to achieve sufficiently high levels of single-phonon non-linearity, an essential ingredient for the reliable isolation of a computational basis. Finally, we will predict competitive performances in the realization of typical digital quantum simulations.

In Chapter 4, we will cross the traditional boundaries of quantum information science, introducing machine learning as a new potential application arena. Like abstract physical theories eventually produced novel insights in computer science, many celebrated protocols in artificial intelligence actually originated from schematic models of neural activity and from the formalization of learning processes. These later evolved into a totally new perspective on software design and led to striking advancements in the way information processing systems interact with the surrounding world: along these lines, the rapidly growing field of Quantum Machine Learning aims at establishing a productive interaction between two currently very active areas of research. Here, we will first give a brief account of the most advanced quantum approaches e.g. to classification tasks, once again with a particular focus on the standard digital and circuit-based model of quantum computation. To provide a fresh merger of quantum information and machine learning, we will then go back to the roots of artificial neural network computation and outline an original procedure for reproduc-

ing, on a quantum architecture, the well known McCulloch-Pitts model of binary-valued artificial neurons. We will describe an optimized data encoding and processing scheme, built upon the class of quantum hypergraph states and showing possible exponential advantages in storage resources over classical equivalents. The protocol, declined as an elementary image recognition problem, will be tested on real IBM Q processors both at the single node level and in a first attempt to realize a fully quantum feed-forward neural network (ffNN) computation. In particular, a proof-of-principle demonstration of a 2-layer and 3-neuron ffNN, run on a 7-qubit quantum register, will be reported and discussed.

A summary of the main achievements, together with an outlook on future works, is left for the final Conclusions.

Digital quantum simulations

A scientific comprehension of the world and its phenomena requires the ability of preparing, observing, registering and interpreting in a coherent framework a series of events. The principle that guides sensible researchers in many aspects of their often wandering activity is in fact best phrased by Richard Feynman's eponymous words "*What I cannot create, I do not understand*". As we live nowadays in a society recognizing the fundamental role of information and information processing as a basis for communication and discovery, it is not surprising that the way scientists, and particularly those acquainted with formal methods of investigation, deal with the recreation of natural facts has also evolved to encompass the computational methods of simulation.

Simulating a natural process means being able to artificially reproduce its properties and its dynamical evolution in time. This is most effectively done in a controlled environment where all external variables are under full control. The very first example of a simulation is, in this extended sense, an accurate mathematical modeling, i.e. a mapping of the information we know about a system of interest onto a certain set of variables and equations, followed by an analytic or numerical solution. The resulting set of mathematical identities, or the computer with the routine aimed at solving them, can then be called a *simulator*. With such a simulator, the experimenter can in principle probe the behavior of the real system under fairly general conditions, initialize its state almost at will, make predictions and test new hypotheses, the only limitations being the validity of the initial modeling and, possibly, the available computational power.

Quantum mechanics, widely recognized as the most successful and so far unchallenged microscopic theory of our universe, currently represents the ultimate frontier of scientific simulation. Models formulated in quantum mechanical terms are often conceptually troublesome and technically demanding. Despite the many methods and tools which have been developed over the last century such as quantum Montecarlo [7], molecular dynamics [8] or tensor networks [9], it has become more and more clear that the distance separating us from the the realm of full many-body quantum simulations will hardly be covered with the sole help of classical computing machines. The most prominent and essential reason lies in the exponential scaling of the size, in terms of time and memory resources, of the numerical problems that researchers typically face when trying to reduce real physical systems to their fundamental constituents. This is in turn a manifestation of the well known quantum mechanical principle associating tensor products of complex vector spaces to systems composed of a large number of parts. The failure of classical simulation methods is especially manifest when strong quantum correlations between subsystems become dominant in determining the properties of the object under investigation, which is indeed the case in many interesting and generally open problems.

Building on the original idea, usually attributed to Feynman [10] with early contributions by Y. I. Manin [11] and P. Benioff [12], that the intimately quantum properties of Nature necessarily require a fully quantum computing machine to be efficiently described, *quantum simulators* have been proposed in the last few decades as the gateway towards a new approach to scientific computing. A quantum simulator is *per se* a physical quantum system which is able, under controlled conditions, to mimic the dynamical behavior of another quantum physical model. In other words, a quantum

simulator arises primarily from the awareness that any computation, and hence any computer, is in itself a physical process [4]. Such an approach immediately leads to the idea of a system which is able to process *quantum* information as if its own internal states obeyed, under a suitable mapping and irrespective of the degree of internal correlations or entanglement, the dynamical laws governing the evolution of a target model. Just like in the early days of classical computing, two complementary viewpoints can be taken on the subject, with different degrees of generality and technical complexity. On one hand, when the focus is put on the physical states of the simulator, namely when the relationship between the controlled and the target system are already evident at the hardware level, we generally speak of *analog* quantum simulators. These devices are tailored to reproduce, on a certain physical platform with its own governing set of quantum dynamical equations, the behavior of a restricted class of other physical models which are also formally governed by the same set of equations, or which can be closely related to it via a simple mapping. The simulator is usually operated by initializing and evolving its state under carefully designed conditions: the observations made on the simulator are then referred to the target model in the very same way a sophisticated toy train can provide hints on its bigger counterpart. Following this route, a plethora of *analog* quantum simulators have been proposed [13–24] and developed [25–36] on several platforms ranging from cold atoms to superconducting circuits, solid state devices and photonic systems. On the other hand, *digital* quantum simulators are programmable and general purpose quantum devices, which promise a larger flexibility on the models to be solved [37–57]. Here the focus is shifted towards the information processing character which ultimately pertains to any physical dynamics and, in this respect, digital quantum simulators are in principle quantum machines satisfying DiVincenzo criteria [58] for universal quantum computation. From now on, we will mainly focus on this second approach, and we will often refer to digital quantum computers as universal quantum simulators (UQS) [37].

In a digital quantum simulation, the evolution of the physical model is mapped, following the usual mathematical formulation of quantum mechanics, onto the effective algebra of quantum registers made of qubits. The quantum time propagator, namely a unitary operation, can then be programmed in digital steps through a sequence of quantum logic gates (i.e. unitary transformations on the qubits) defining a quantum circuit [5]. In the following, we will provide specific examples and general recipes both for the mapping of the target system and for the translation of evolution operators into quantum register operations. We will in particular concentrate on the class of quantum spin models, which besides being extremely interesting on their own right, usually constitute the ideal formal conjunction between general quantum mechanical models and their corresponding representation in terms of qubits. Spin models are in this sense the key towards the quantum simulation of a large class of manybody quantum models, such as interacting fermionic systems, that are known to be classically intractable [59]. Paradigmatic examples are the Hubbard model in condensed matter [48, 60], or the Schwinger model in lattice gauge field theory [52, 61, 62]. In particular, we will emphasize the role of specific quantities that are known to be difficult to compute but extremely important in the description of the dynamical response [63] of manybody systems, such as quantum correlations. Despite being in principle much more powerful, UQS are typically difficult to realize in practice compared to analog simulators, mainly due to the well known stringent requirements for general purpose quantum computation [58]. In this respect, it is worth mentioning that hybrid digital-analog quantum simulators have also been proposed, aimed at combining the easier scalability of analog approaches with the intrinsic universality of digital quantum simulations [64]. This paradigm is hailed as a promising route leading to universal digital-analog quantum computation.

Several general purpose [6, 65–67] and more category- or hardware-specific [68–73] accounts of the development in the field of quantum simulators can be found in the literature. In the following sections, we will start by giving a detailed description of the theoretical foundations of digital quantum simulations, commenting on the related mathematical techniques and with a clear focus on the tools that are most often required in practical cases. This will provide the necessary background to support the description of the original results presented in the following chapters. We will then describe and comment the recent and future development of the field in terms of

both algorithmic procedures and experimental results. We will pay special attention to near-term realizations of digital quantum computers and simulators, and particularly to those technological platforms which are currently leading the way, namely coherently manipulated trapped ions [74, 75] and superconducting circuits working at microwave frequencies [76–78]. We will closely follow the fast pace of advancement which, in recent years, has made programmable devices available even outside research laboratories, attracting widespread interest. Indeed, established technological companies such as IBM and Google, startup ventures (Rigetti Computing and IonQ to name a few) and public institutions [79] all have designed and deployed consistent strategies in search of the long-sought goal of “quantum advantage”. Such efforts hold promise for breaking the barriers currently limiting the simulations of complex many-body physics with classical computing machines. Without entering into rigorous definitions here [80–82], we can identify such threshold close to the size of 50 to 60 fully operational qubits. Indeed, by reasoning in orders of magnitude and by assuming that ~ 8 bytes are needed to store a complex number in single-precision, a 50 qubit quantum register would in general be able to manipulate around $8 \times 2^N \sim 9 \cdot 10^{15}$ bytes of information, corresponding to roughly 9 Pb. This in turn approximates the typical amount of random access memory in state-of-the-art supercomputers [80, 81]. Quantum advantage for scientific applications is largely believed to be within reach already in the current era of Noisy Intermediate Scale Quantum (NISQ) devices [6], and it represents the landmark which would unequivocally certify the maturity of the field and probably its potential commercial value. Towards this direction, relevant claims of a 53-qubit superconducting quantum hardware outperforming even the most powerful supercomputer currently available in the completion of a specific algorithmic task have recently been reported [83]. While certainly setting an important technological milestone, such device has not yet been employed to run a programmable quantum simulation of a physical model with certified advantage, nor it has been used so far for the demonstration of practical use cases. Far-reaching consequences may of course be expected if fully fault tolerant and scalable quantum hardware will effectively become available [60, 84–86], in which $N > 100$ logical qubits should be complemented with a much larger number of auxiliary quantum bits aimed at correcting noise induced errors. However, it is more difficult to foresee now a timescale for the practical realization of this paradigm.

On a parallel sight, while the main object of this chapter is restricted to quantum simulations of physical models, it is nevertheless worth mentioning already at this stage that the same techniques could in principle be applied to more general computational tasks. Some complex problems in fields outside the physical sciences, such as optimization, stock market pricing [87, 88] and machine learning [89] are indeed known to have close relationships with many mathematical models in physics or engineering, and might therefore benefit from speedup advantages over classical computers.

1.1 Theory of quantum simulations

When a physical theory is designed to describe the evolution in time of a system, the mathematical description is typically formulated in the language of differential equations. Their solution is at the heart of most simulation protocols nowadays, from molecular dynamics to aircraft design. A very common situation is, for example, a linear set of equations such as

$$\frac{d\vec{x}}{dt} = M\vec{x} \quad (1.1)$$

where M is a matrix and \vec{x} represents a vector of dynamical variables. Given an initial condition $\vec{x}(0)$, the formal solution to the above equation is simply

$$\vec{x}(t) = e^{Mt}\vec{x}(0) \quad (1.2)$$

Implementing such a solution on a computer routine gives a useful tool to fully solve the system dynamics, provided that the size of the numerical problem is within reach of the available computational resources. In quantum mechanics, the corresponding paradigmatic example is the Schrödinger equation (here and in the following, we will assume $\hbar = 1$)

$$\frac{d|\Psi\rangle}{dt} = -i\mathcal{H}|\Psi\rangle \quad (1.3)$$

where \mathcal{H} is known as the Hamiltonian operator and the associated complex linear space is the well known Hilbert space of wavefunctions. The Schrödinger equation can in principle be fully solved by computing the unitary time-evolution operator $U(t) = e^{-i\mathcal{H}t}$. Once the latter is known, any initial condition can be evolved linearly as

$$|\Psi(t)\rangle = U(t) |\Psi(0)\rangle \quad (1.4)$$

From the few examples above, we can already see that matrix exponentiation is an ubiquitous numerical task in simulation scenarios, and crucially in the field of quantum mechanical systems. On classical computers, matrix exponentiation is a notably difficult problem whose computational complexity is believed to scale at least polynomially with the size of the matrix [90]. When combined with the exponential growth of the linear dimensions associated to a composite quantum mechanical system with the number of sub-systems, this in turn leads in general to an exponential demand of time and memory resources for manybody quantum simulations. In the next section we will discuss how quantum computing devices might be able to overcome such limitations in many cases of practical interest.

1.1.1 From Feynman to Lloyd

In 1982, Richard Feynman conjectured that using a controllable quantum mechanical system as a computing resource, instead of a classical object, would provide significant advantages in the simulation of quantum systems [10]. Indeed, just about fifteen years later, in 1996, Seth Lloyd proved that idea to be essentially correct [37], with the sole limitation that the systems to be simulated only carry local interactions between their constituent subsystems. We will thus concentrate on system Hamiltonians of the form

$$H = \sum_l^L H_l \quad (1.5)$$

where H_l acts locally only on a portion of the total system.

Given a certain Hamiltonian that models the physical system under investigation, H , the problem of computing the corresponding time evolution operator $U(t) = \exp(-iHt)$ is equivalent, from a quantum informational perspective, to the task of realizing a well defined unitary transformation. As it is well known in the field [5], a quantum computing device endowed with a universal set of quantum gates is by definition able to perform any arbitrary unitary transformation, albeit not necessarily in an efficient number of elementary operations. What Lloyd actually proved is that universal quantum computers can calculate $U(t)$ efficiently (i.e. with polynomial time and memory resources in the size of the target system) whenever H is a sum of local terms. The proof of this result relies on the following two fundamental facts: first, in the circuitual model for universal quantum computation a general unitary transformations can always be implemented by successively performing elementary unitary operations (quantum gates), and appending one unitary U_A after another U_B in the circuit results in a total unitary which mathematically is the product $U_A U_B$ being applied to the state of the qubit register. Second, any unitary operation U acting on N qubits can be implemented with $O(2^{2N})$ elementary operations [5, 91]. Here we recall that the dimension of the Hilbert space associated to N qubits is $d = 2^N$. Suppose now that we are given a Hamiltonian which is a sum of local terms, as in Eq. (1.5), with say $L \propto p \cdot N$, where p measures some degree of locality¹ and N is the total number of qubits required to encode the computation. According to the rules above, computing directly $U(t) = \exp(-iHt)$ in general requires $O(2^{2N})$ operations, and is therefore exponentially inefficient. However, let us call m_l the dimension of the subsystem over which the action of each H_l is restricted. We can assume that $m_l \ll 2^N$, since typical local terms only involve few-body interactions. In this case, the unitary $U_l(t) = \exp(-iH_l t)$ can be computed with $O(m_l^2)$ operations. The overall product

$$\tilde{U} = \prod_l U_l(t) \quad (1.6)$$

¹For example, in a lattice of N nodes with pairwise interactions and p nearest neighbors we have $L = pN/2$.

can therefore be obtained on a universal quantum computer by juxtaposing the circuit implementations of the single $U_l(t)$ unitary matrices and takes at most $O(Lm_{max}^2)$ elementary operations, where $m_{max} = \max_l m_l$. The final step of the reasoning lies in the following mathematical identity, which is known as the Suzuki-Trotter (ST) decomposition:

$$e^{-i\sum_l H_l t} = \lim_{n \rightarrow \infty} \left(\prod_l e^{-iH_l t/n} \right)^n \quad (1.7)$$

If all the H_l operators commute, namely

$$[H_l, H_{l'}] = 0 \forall l, l' \quad (1.8)$$

the ST identity is exact already for $n = 1$. However, in general the product of local time evolution operators will not be exactly equal to the total target unitary $U(t) = \exp(-iHt)$, but it can be shown that $\forall n$

$$U(t) = e^{-i\sum_l H_l t} = \left(\prod_l e^{-iH_l t/n} \right)^n + O\left(\frac{t^2}{n}\right) \quad (1.9)$$

The equation above means that we can approximate arbitrarily well the desired quantum mechanical time propagator by repeating n times the sequence of gates corresponding to the product of local terms for shorter time steps t/n . All in all, we were able to break our original problem into smaller pieces, $e^{-iH_l t/n}$, which can now be implemented efficiently using only a limited set of elementary gates and give the correct answer up to an arbitrarily small digital error $O(t^2/n)$. Indeed, for any $\epsilon > 0$ and t , there exists a n_ϵ such that $U(t)$ can be computed within an approximation ϵ in at most $n_\epsilon L m_{max}^2$ operations. This is polynomial in N whenever $L = \text{poly}(N)$, as for example in the case of nearest neighbors interactions.

1.1.2 Quantum simulations cookbook

A universal quantum computer will be described, from now on, as a qubit-based digital quantum device obeying the algebra of Pauli matrices and operating with a universal set of quantum gates [5]. Given the results presented in the previous section, the problem of quantum simulation can then be formulated and solved on such a machine in a handful of simple steps.

First, a model Hamiltonian of interest \mathcal{H} must be defined. As for any sensible physical description, \mathcal{H} should contain all the dynamical information characterizing the physical quantum system under investigation. The most appropriate set of variables and operators will appear in the mathematical structure of the Hamiltonian, such as for example frequencies and couplings, spin matrices, fermionic/bosonic creation and annihilation operators or lattice-discretized quantum fields.

Second, the target Hamiltonian \mathcal{H} must be mapped onto its equivalent representation on the qubit Pauli algebra

$$\mathcal{H} \rightarrow H(\{\sigma_\alpha\}) \quad (1.10)$$

This step requires a suitable encoding of the degrees of freedom of the target system into a number N of qubits. All the relevant quantum mechanical states and operators must be rephrased in terms of computational basis states and Pauli matrices $\{\sigma_\alpha\}$ acting on them, thus resulting in mapped Hamiltonian H . We recall here that the Pauli algebra of qubits is characterized by the following well known set of $SU(2)$ matrices

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (1.11)$$

satisfying the following commutation and anti-commutation rules

$$[\sigma_\alpha, \sigma_\beta] = 2i\epsilon_{\alpha\beta\gamma}\sigma_\gamma, \quad \{\sigma_\alpha, \sigma_\beta\} = 2\delta_{\alpha\beta}\mathbb{I} \quad (1.12)$$

where $\alpha, \beta, \gamma \in \{x, y, z\}$, $\epsilon_{\alpha\beta\gamma}$ is the Levi-Civita tensor, $\delta_{\alpha\beta}$ is the Kronecker delta and \mathbb{I} is the identity matrix. The Pauli algebra is at the heart of the physical description of spin-1/2 quantum

systems and therefore the mapped Hamiltonian H will in general correspond to a model of interacting spin-1/2 operators, such as for example the Heisenberg or Ising chain. It is easy to understand at this stage the reason why spin systems represent the ideal contact point between physical and computational quantum problems. Despite being intrinsically more involved, effective mappings are also known for a large variety of interesting cases, ranging from spin $S > 1/2$ [47, 54, 92] to fermionic and fermionic-bosonic systems [38, 47–49, 54, 57, 64, 93–95] usually through the well known Jordan-Wigner transformation [40, 47, 96, 97]. Other recent examples include lattice models related to gauge theories [61, 62] and even quantum gravity studies [55]. Regardless of the details of the original \mathcal{H} , the quantum simulation will be efficient, in Lloyd’s sense, whenever the corresponding H is the sum of local terms. This is usually not a limitation in many practical cases, as most physical processes are inherently local in nature.

In the third step, assuming the target Hamiltonian is mapped onto a sum of local contributions

$$H = \sum_l H_l \quad (1.13)$$

one must check whether $[H_l, H_{l'}] = 0 \forall l, l'$. If that is the case, then

$$e^{-iHt} = \prod_l e^{-iH_l t} \quad (1.14)$$

with no error. Otherwise, an appropriate number n of ST steps (sometimes referred to as Trotter steps) must be chosen according to the required degree of precision, in such a way that, by applying Equations (1.7) and (1.9)

$$e^{-iHt} \simeq \left(\prod_l e^{-iH_l t/n} \right)^n \quad (1.15)$$

This application of the ST formula is sometimes called *Trotterization* in quantum simulations jargon, and the error which arises from the above approximation formula is also named *digitalization* or *digital* error. In pictorial terms, one might somehow compare this situation with the well known analog-to-digital conversion tasks performed on classical computers, where a finite set of discrete operations accessible to the machine must be used to approximate a continuum of possible signals. We will come back to the properties and limitations set by digitalization errors in a subsequent section.

Fourth, each local unitary $e^{-iH_l t}$ (or $e^{-iH_l t/n}$) must be translated into a sequence of quantum gates. This is always possible in at most $O(m_l^2)$ operations and with any universal set of single- and two-qubit operations available on a general purpose quantum computer [5]. However, no unique scheme or library exists in general to carry out such translation, as different universal set of operations are in principle all equivalent and specific choices are usually dictated by the processor architecture or by technical efficiency considerations. Some examples and common “dictionaries” will be discussed later in this chapter. Once the individual sequences corresponding to the factors in the ST decomposition are known, the total quantum circuit encoding the time evolution will simply be the juxtaposition of all of them, repeated in general n times.

Finally, initial state preparation must be included at the beginning of the circuit, and an appropriate set of measurements must be appended at the end to recover expectation values of the relevant observable quantities on the evolved quantum state. Notice that these are by no means trivial tasks, since general quantum states and measurements require the ability of preparing and observing highly correlated states and properties.

While the points outlined above represent a quite general set of instructions towards the design of a quantum simulation algorithm, moving to practical implementations of such scheme usually requires a number of algorithmic and quantum computational tools, which must also be adapted to the specific hardware platform on which the computation is to be performed. We will provide specific examples and useful formal techniques in the following sections. The general philosophy behind a digital quantum simulation is also summarized in Fig. 1.1. Once more, we stress that the class of computational tasks which can be treated with the methods of quantum simulations is not at all limited to actual physical systems, but extends towards any problem that is expressible in the form of Hamiltonian quantum dynamics.

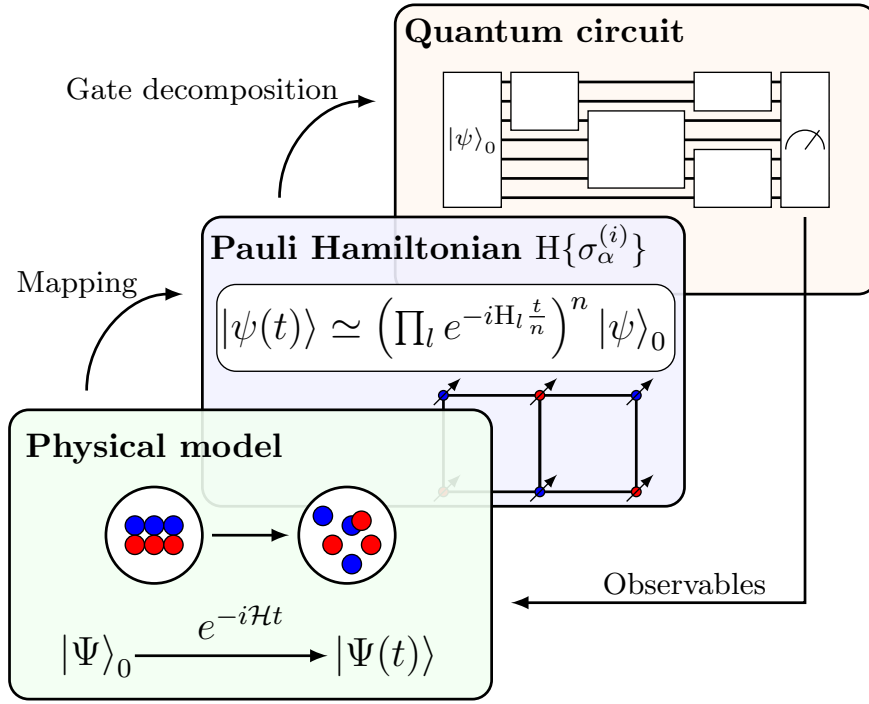


Figure 1.1: Conceptual illustration of a digital quantum simulator. A physical Hamiltonian \mathcal{H} determines the quantum state evolution $\Psi(t)$, which can be approximated to arbitrary precision by mapping the given model on a spin-type algebra and slicing the time evolution according to the Trotter-Suzuki formula. The resulting sequence of unitary operations is then put in the form of a quantum circuit to be directly run on a quantum computing hardware. Finally, observations made on the approximated output state of the quantum register, $\psi(t)$, are mapped back to the original physical space.

1.1.3 Spin mapping and circuit decompositions

Among the steps that must be undertaken in order to practically design and realize a digital quantum simulation, the translation of unitary operators into elementary quantum gates is the one that is most typically hardware-dependent. It is also critical in terms of results and performance, particularly in the present era of NISQ prototypes, where the correlation between hardware properties and target features is stronger. In this section, after introducing the general form of typical mapped Hamiltonians \mathcal{H} , we will review the most common choices for universal sets of single- and two-qubit gates and we will relate them to the technological platforms where they are employed. We will then present, for paradigmatic spin-type Hamiltonian terms, the corresponding possible quantum gate sequences: in this way, we are going to build a fundamental library that will constitute the stepping stone towards more specific applications. We recall, as already mentioned before, that the choice of spin models does not limit the generality of the discussion: indeed, they formally obey the abstract Pauli algebra which is the common logical basis of all digital quantum registers.

In Section 1.1.2 we have introduced the general idea that, for all practical purposes, most problems in digital quantum simulations are ultimately brought in the form of a spin-1/2 Hamiltonian which can be directly mapped onto a N -qubit quantum register. As an example [73, 93], consider the following Fermi-Hubbard minimal model describing two lattice sites $i \in \{1, 2\}$ and two fermions with opposite spin directions $s \in \{\uparrow, \downarrow\}$:

$$\mathcal{H} = -V \left(b_{1,\downarrow}^\dagger b_{2,\downarrow} + b_{1,\uparrow}^\dagger b_{2,\uparrow} + h.c. \right) + U \left(b_{1,\downarrow}^\dagger b_{1,\downarrow} b_{1,\uparrow}^\dagger b_{1,\uparrow} + b_{2,\downarrow}^\dagger b_{2,\downarrow} b_{2,\uparrow}^\dagger b_{2,\uparrow} \right) \quad (1.16)$$

where $h.c.$ means the Hermitian conjugate and the operator $b_{i,s}^\dagger$ ($b_{i,s}$) creates (annihilates) a fermion with spin s on site i . In \mathcal{H} , the coefficients V and U are hopping and on-site repulsion energies.

Fermionic operators obey the canonical anticommutation rules

$$\{b_{i,s}, b_{i',s'}^\dagger\} = \delta_{ii',ss'} \mathbb{I} \quad \{b_{i,s}, b_{i',s'}\} = 0 \quad (1.17)$$

The spin mapping can be obtained by applying the Jordan-Wigner transformation [38, 40, 96, 97], which in this case takes the explicit form

$$\begin{aligned} b_{1,\downarrow}^\dagger &= \mathbb{I} \otimes \mathbb{I} \otimes \mathbb{I} \otimes \sigma^+ \\ b_{2,\downarrow}^\dagger &= \mathbb{I} \otimes \mathbb{I} \otimes \sigma^+ \otimes \sigma_z \\ b_{1,\uparrow}^\dagger &= \mathbb{I} \otimes \sigma^+ \otimes \sigma_z \otimes \sigma_z \\ b_{2,\uparrow}^\dagger &= \sigma^+ \otimes \sigma_z \otimes \sigma_z \otimes \sigma_z \end{aligned} \quad (1.18)$$

Here $2\sigma^+ = \sigma_x + i\sigma_y$ and the anticommutation relations are preserved due to the properties of Pauli matrices. The resulting mapped Hamiltonian reads

$$\begin{aligned} \mathbb{H} &= \frac{V}{2} (\sigma_x \otimes \sigma_x \otimes \mathbb{I} \otimes \mathbb{I} + \sigma_y \otimes \sigma_y \otimes \mathbb{I} \otimes \mathbb{I} + \mathbb{I} \otimes \mathbb{I} \otimes \sigma_x \otimes \sigma_x + \mathbb{I} \otimes \mathbb{I} \otimes \sigma_y \otimes \sigma_y) \\ &+ \frac{U}{4} (\sigma_z \otimes \mathbb{I} \otimes \mathbb{I} \otimes \sigma_z + \mathbb{I} \otimes \mathbb{I} \otimes \mathbb{I} \otimes \sigma_z + \sigma_z \otimes \mathbb{I} \otimes \mathbb{I} \otimes \mathbb{I}) \\ &+ \frac{U}{4} (\mathbb{I} \otimes \sigma_z \otimes \sigma_z \otimes \mathbb{I} + \mathbb{I} \otimes \sigma_z \otimes \mathbb{I} \otimes \mathbb{I} + \mathbb{I} \otimes \mathbb{I} \otimes \sigma_z \otimes \mathbb{I}) \end{aligned} \quad (1.19)$$

or, equivalently,

$$\begin{aligned} \mathbb{H} &= \frac{V}{2} \left(\sigma_x^{(1)} \sigma_x^{(2)} + \sigma_y^{(1)} \sigma_y^{(2)} + \sigma_x^{(3)} \sigma_x^{(4)} + \sigma_y^{(3)} \sigma_y^{(4)} \right) \\ &+ \frac{U}{4} \left(\sigma_z^{(1)} \sigma_z^{(4)} + \sigma_z^{(1)} + \sigma_z^{(4)} + \sigma_z^{(2)} \sigma_z^{(3)} + \sigma_z^{(2)} + \sigma_z^{(3)} \right) \end{aligned} \quad (1.20)$$

where we indicate with $\sigma_\alpha^{(k)}$ the operator σ_α acting on the k -th component of the tensor product. As \mathbb{H} is now the sum of local Pauli terms, the corresponding time evolution operator could be simulated on a 4-qubit quantum register, as it will be clear in the following. It is not difficult to abstract from this elementary example and to write down the typical form for the generator of time evolution in a N -qubit digital quantum simulation:

$$\mathbb{H} = \sum_{\substack{i=1 \\ \alpha=x,y,z}}^N h_{\alpha,i}^{(1)} \sigma_\alpha^{(i)} + \sum_{\substack{i,j=1 \\ \alpha,\beta=x,y,z}}^N h_{\alpha\beta,ij}^{(2)} \sigma_\alpha^{(i)} \sigma_\beta^{(j)} + \dots \quad (1.21)$$

Only single- and two-spin terms are explicitly shown here, as the time evolution induced by any other manybody term can in principle be reduced to these ones in a direct application of the well known universality of single- and two-qubit operations [5]. However, it is worth mentioning that in some specific architectures where many-body interactions are natively implementable (e.g. in trapped ions quantum processors) the class of elementary operations can be extended without affecting neither the universality nor the efficiency of the quantum simulation methods, provided that the overall structure of \mathbb{H} retains a local nature. We should also notice at this stage that $\mathcal{H} \mapsto \mathbb{H}(\{\sigma_\alpha\})$ mapping procedures may generate non-local or many-body interactions even from apparently simple target Hamiltonians \mathcal{H} : for example, the application of the Jordan-Wigner transformation in more than one dimension or with long-range couplings leads to multi-spin interactions in the resulting \mathbb{H} . Indeed, Eq. (1.18) must in general be replaced by $b_j^\dagger = \left(\prod_{l=1}^{j-1} -\sigma_z^{(l)} \right) \sigma_+^{(j)}$ [40], where we have introduced for simplicity a single label $j = \{i, s\}$ for fermionic modes addressing both site and spin variables. The resulting spin Hamiltonian contains several terms, depending on the number of occupied states existing between each pair of interacting fermions after lattice sites have been properly ordered [54].

Before moving to technical discussions, we can now use Eq. (1.21) to set up the notation and terminology for some of the renown spin models that we will most often deal with in the rest of this work. Among these, the Heisenberg model corresponds to a Hamiltonian of the form

$$H_{\text{Heis}} = J \sum_{\langle i,j \rangle} \left(\sigma_x^{(i)} \sigma_x^{(j)} + \sigma_y^{(i)} \sigma_y^{(j)} + \sigma_z^{(i)} \sigma_z^{(j)} \right), \quad (1.22)$$

A generalization to anisotropic couplings gives the XYZ model

$$H_{\text{XYZ}} = \sum_{\langle i,j \rangle} \left(J_{xx} \sigma_x^{(i)} \sigma_x^{(j)} + J_{yy} \sigma_y^{(i)} \sigma_y^{(j)} + J_{zz} \sigma_z^{(i)} \sigma_z^{(j)} \right), \quad (1.23)$$

which reduces to the so called XY model if $J_{zz} = 0$. In all cases, a transverse magnetic field can be added via single qubit terms, such as in the Transverse Field Ising model (TIM)

$$H_{\text{TIM}} = \sum_i h_i \sigma_x^{(i)} + \sum_{\langle i,j \rangle} J_{zz} \sigma_z^{(i)} \sigma_z^{(j)} \quad (1.24)$$

Here $\langle i, j \rangle$ denote nearest neighbors spin pairs.

Universal sets of quantum gates on NISQ processors

Any given hardware architecture comes with a native set of operations which are implemented in practice according to the physical characteristics of the device. The platform has the potential for implementing universal quantum computation and simulation if and only if, among all other criteria [58], such native set of operations contains a universal set in the usual quantum computing sense. Several of such universal sets of single- and two-qubit gates are known [5], all in principle equally powerful. If the requirement is satisfied, any target unitary evolution can be translated into a combination of the native operations without unnecessary overhead. Processors based on different technological platforms may also exhibit distinct qubit-qubit inter-connectivity and other more technical constraints such as limitations in gate directionality. These features do not usually pose hard limitations to the computational power of the platform, since they can always be compensated via, e.g., SWAP operations. However, they may result in some overhead in the total length of the simulations, even up to the point that the results of certain computations are significantly degraded due to the still limited resilience of NISQ devices to noise, cross-talks and gate errors. As a consequence, it is certainly the case that some NISQ platforms are more suitable for the simulation of certain physical models (e.g. trapped ions, featuring built-in all-to-all connectivity, can more easily simulate long-range interactions), even though in principle any full-fledged digital quantum simulator could reproduce arbitrary quantum dynamics. A fair comparison of performances between competing platforms requires all these factors to be taken into account [98, 99].

Single qubit control is one of the most basic requirements for general-purpose quantum computation. All currently proposed and realized quantum computing platforms allow addressing single qubits with tailored control pulses to perform single qubit gates. The most general single qubit SU(2) operation has the form

$$U(\theta, \phi, \lambda) = \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i(\lambda+\phi)} \cos(\theta/2) \end{pmatrix} \quad (1.25)$$

and can be obtained, for example, by combining well known single qubit quantum gates such as the Hadamard gate²

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (1.26)$$

and the phase gate

$$\Phi(\delta) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\delta} \end{pmatrix} \quad (1.27)$$

²There is often a conflict of notation between Hadamard gates and Hamiltonians, for both are customarily indicated by ‘‘H’’. We will not unnecessarily complicate the notation as long as the meaning is clear from the context.

Indeed, the following identity holds:

$$U(\theta, \phi, \lambda) = e^{-i\theta/2} \Phi\left(\frac{\pi}{2} + \phi\right) H\Phi(\theta) H\Phi\left(-\frac{\pi}{2} + \lambda\right) \quad (1.28)$$

Rotations around the coordinate axes

$$R_\alpha(\theta) = \exp\left(-i\frac{\theta}{2}\sigma_\alpha\right) \quad \alpha = x, y, z \quad (1.29)$$

can be implemented, up to global phase factors, by choosing particular parameters in $U(\theta, \phi, \lambda)$. For example, $R_z(\lambda) = e^{-i\lambda/2}\Phi(\lambda) = e^{-i\lambda/2}U(0, 0, \lambda)$, $R_x(\theta) = U(\theta, -\pi/2, \pi/2)$ and $R_y(\theta) = U(\theta, 0, 0)$. Vice-versa, any platform capable of implementing single-qubit rotations around the coordinate axes can in principle realize an arbitrary $U(\theta, \phi, \lambda)$ via the identity

$$U(\theta, \phi, \lambda) = R_z(\phi)R_x(\theta)R_z(\lambda) \quad (1.30)$$

Finally, it is interesting to mention that the mathematical relation

$$U(\theta, \phi, \lambda) = R_z(\phi - \pi/2)R_x(\pi/2)R_z(\pi - \theta)R_x(\pi/2)R_z(\lambda - \pi/2) \quad (1.31)$$

justifies a particularly effective experimental technique in which arbitrary single-qubit operations are performed by using only fixed-phase $R_x(\pi/2)$ gates, thus requiring just a single or few external pulse calibrations, and error-free *virtual* R_z gates which are essentially changes of reference frame in the control software [100]. Other formal decompositions of general $SU(2)$ transformations, as well as approximate results employing only a finite set of fixed-phase single qubit operations instead of continuous-valued ones, are also known [5, 101].

All universal sets of quantum gates require at least a two-qubit fully entangling operation. The exact way this requirement is satisfied in real processors strongly depends on the underlying architecture and might differ even for competing implementations based on the same fundamental technology. For example, superconducting qubits connected via cross-resonance interactions [102–105] often display a native universal set

$$\mathcal{S}_1 = \{R_\alpha(\theta), \text{CNOT}\} \quad (1.32)$$

containing arbitrary single qubit rotations and the two-qubit CNOT gate

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (1.33)$$

A parametric alternative to the CNOT which is usually very flexible and well suited for digital quantum simulations is expressed in the form of a so called XY interaction [53, 92, 106, 107]

$$U_{XY}(\delta) = e^{-i\delta(\sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y)} \quad (1.34)$$

In matrix notation, this corresponds to

$$U_{XY}(\delta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(2\delta) & -i \sin(2\delta) & 0 \\ 0 & -i \sin(2\delta) & \cos(2\delta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1.35)$$

which for $\delta = -\pi/8$ gives the $\sqrt{i\text{SWAP}}$ gate

$$U_{\sqrt{i\text{SWAP}}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1/\sqrt{2} & i/\sqrt{2} & 0 \\ 0 & i/\sqrt{2} & 1/\sqrt{2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1.36)$$

In contrast with the previous case, at the hardware level the XY interaction, which is equivalently generated by an exchange Hamiltonian of the form

$$H_{\text{hardware}} \propto (\sigma^+ \otimes \sigma^- + \sigma^- \otimes \sigma^+) \quad (1.37)$$

arises naturally, e.g., in many direct coupling configurations of superconducting circuits [72, 108, 109]. Here, we will call \mathcal{S}_2 the universal set

$$\mathcal{S}_2 = \{R_\alpha(\theta), U_{XY}(\delta)\} \quad (1.38)$$

Finally, we define $\mathcal{S}_3 = \{R_\alpha(\theta), C\Phi(\delta)\}$ as the universal set of quantum gates containing all single qubit rotations and the controlled-phase gate

$$C\Phi(\delta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{i\delta} \end{pmatrix} \quad (1.39)$$

The latter is often natively implemented on superconducting platforms with state dependent frequency shifts [60, 93, 110–112], and it is closely related to the Ising interaction generated by $H_{\text{Ising}} \propto \sigma_z \otimes \sigma_z$ [113].

Shifting our focus to quantum processors based on trapped ions technology [75, 114, 115], the most widely used fundamental set of operations, which we will call \mathcal{S}_4 , includes individual single qubit z -rotations

$$T_1^{(j)}(\theta) = e^{-i\theta\sigma_z^{(j)}}, \quad (1.40)$$

collective non-entangling operations

$$T_2(\theta) = e^{-i\theta\sum_j\sigma_z^{(j)}}, \quad T_3(\theta, \phi) = e^{-i\theta\sum_j\sigma_\phi^{(j)}} \quad (1.41)$$

where $\sigma_\phi = \cos\phi\sigma_x + \sin\phi\sigma_y$, and Mølmer-Sørensen collective entangling gates [116–118]

$$T_4(\theta, \phi) = e^{-i\theta\sum_{i<j}\sigma_\phi^{(i)}\sigma_\phi^{(j)}} \quad (1.42)$$

Any subset of qubits can in principle be addressed with the collective gates via optical laser manipulation and transverse normal vibrational modes, while leaving the others untouched. This collective character of trapped ions quantum gates is best exploited for the quantum simulation of long range and multiple-body interactions, as it will be detailed in Sec 1.2.1.

Library of quantum gates sequences for quantum simulations

Once the hardware-native fundamental set of operations is known, any given quantum simulation is most effectively digitized by making use of the single- and two-qubit operations which are accessible to the quantum processor. Every spin-1/2 appearing in the Hamiltonian of the form shown in Eq. (1.21) is assigned to a qubit in the quantum register, and the corresponding unitary manipulations are then applied.

Starting from single-spin terms

$$H_1^{(i)} = \sum_{\alpha=x,y,z} h_{\alpha,i}^{(1)} \sigma_\alpha^{(i)}, \quad (1.43)$$

these induce a time evolution

$$U_1^{(i)}(t) = e^{-iH_1^{(i)}t} \quad (1.44)$$

corresponding to a rotation of the Bloch vector of the target qubit. As shown in the previous section, this can always be expressed in the $U(\theta, \phi, \lambda)$ form, and therefore as a combination of rotations around the coordinate axes or alternatively of e.g. Hadamard and phase gates.

Two-body spin-spin interaction terms can in general be written as

$$H_{\alpha\beta}^{(i,j)} = h_{\alpha\beta,ij}^{(2)} \sigma_{\alpha}^{(i)} \otimes \sigma_{\beta}^{(j)} \quad (1.45)$$

and the corresponding local time evolution operator reads

$$U_{\alpha\beta}^{(i,j)}(t) = e^{-iH_{\alpha\beta}^{(i,j)}t} = e^{-i\delta\sigma_{\alpha}^{(i)} \otimes \sigma_{\beta}^{(j)}} \quad (1.46)$$

where we take δ as a dimensionless phase factor. The decomposition of $U_{\alpha\beta}^{(i,j)}(t)$ terms into elementary quantum gates varies depending on the chosen set of native operations. For example, let us define the unitary operation $ZZ(\delta)$ as

$$ZZ(\delta) = e^{-i\delta\sigma_z \otimes \sigma_z} \quad (1.47)$$

Within the universal set \mathcal{S}_1 , $ZZ(\delta)$ can be implemented according to the following quantum circuit

$$e^{-i\delta\sigma_z \otimes \sigma_z} = \begin{array}{c} \bullet \text{---} \text{---} \bullet \\ | \quad | \\ \oplus \text{---} \boxed{R_z(2\delta)} \text{---} \oplus \end{array} \quad (1.48)$$

where we employed the usual notation for CNOT gates

$$\text{CNOT} = \begin{array}{c} \bullet \\ | \\ \oplus \end{array} \quad (1.49)$$

Similar constructions, combined with single qubit rotations, give access to all the $U_{\alpha\beta}^{(i,j)}(t)$ transformations. Indeed, by recalling the following identities

$$\begin{aligned} R_y\left(\frac{\pi}{2}\right) \sigma_z R_y\left(-\frac{\pi}{2}\right) &= \sigma_x \\ R_x\left(\frac{\pi}{2}\right) \sigma_z R_x\left(-\frac{\pi}{2}\right) &= -\sigma_y \end{aligned} \quad (1.50)$$

which essentially correspond to rotations of the qubit reference frame, it is straightforward to verify that the two-qubit operation

$$YY(\delta) = e^{-i\delta\sigma_y \otimes \sigma_y} \quad (1.51)$$

can be obtained in with the following sequence of gates

$$e^{-i\delta\sigma_y \otimes \sigma_y} = \begin{array}{c} \boxed{R_x(\pi/2)} \text{---} \bullet \text{---} \bullet \text{---} \boxed{R_x(-\pi/2)} \\ | \quad | \\ \boxed{R_x(\pi/2)} \text{---} \oplus \text{---} \boxed{R_z(2\delta)} \text{---} \oplus \text{---} \boxed{R_x(-\pi/2)} \end{array} \quad (1.52)$$

and, for example,

$$e^{-i\delta\sigma_x \otimes \sigma_z} = \begin{array}{c} \boxed{R_y(-\pi/2)} \text{---} \bullet \text{---} \bullet \text{---} \boxed{R_y(\pi/2)} \\ | \quad | \\ \oplus \text{---} \boxed{R_z(2\delta)} \text{---} \oplus \end{array} \quad (1.53)$$

When using the universal set \mathcal{S}_2 , it is easier to take as the fundamental building block the unitary transformation

$$XX(\delta) = e^{-i\delta\sigma_x \otimes \sigma_x} \quad (1.54)$$

which corresponds to the following quantum circuit

$$e^{-i\delta\sigma_x \otimes \sigma_x} = \begin{array}{c} \boxed{U_{XY}(\delta/2)} \text{---} \boxed{R_x(\pi)} \text{---} \boxed{U_{XY}(\delta/2)} \text{---} \boxed{R_x(-\pi)} \\ | \quad | \quad | \\ \text{---} \text{---} \text{---} \end{array} \quad (1.55)$$

The decomposition above is readily understood by considering that

$$\mathbf{R}_x(-\pi)\sigma_y\mathbf{R}_x(\pi) = -\sigma_y \quad (1.56)$$

and using the identity

$$e^{-i\delta\sigma_x\otimes\sigma_x} = e^{-i\frac{\delta}{2}(\sigma_x\otimes\sigma_x-\sigma_y\otimes\sigma_y)}e^{-i\frac{\delta}{2}(\sigma_x\otimes\sigma_x+\sigma_y\otimes\sigma_y)} \quad (1.57)$$

All other unitary evolution terms generated by $\sigma_\alpha \otimes \sigma_\beta$ can again be reduced to $\mathbf{XX}(\delta)$ and to the quantum circuit in Eq. (1.55) with single-qubit changes of reference frame.

In view of the already mentioned close relationship of $\mathbf{C}\Phi(\delta)$ gates with Ising-like evolution, it is not surprising that the $\mathbf{ZZ}(\delta)$ building block can be obtained within \mathcal{S}_3 directly from a single $\mathbf{C}\Phi(\delta)$ with single qubit corrections and apart from an overall phase:

$$e^{i(\delta/4)}\mathbf{ZZ}(\delta/4) = \begin{array}{c} \boxed{\Phi(-\delta/2)} \text{---} \bullet \text{---} \\ | \\ \text{---} \boxed{\Phi(\delta)} \text{---} \boxed{\Phi(-\delta/2)} \text{---} \end{array} \quad (1.58)$$

Since in real experimental setups the achievable phases δ in a single $\mathbf{C}\Phi(\delta)$ gate might be limited due to hardware constraints, for example in a practical range $\delta = 0.5\text{--}4.0$ rads [93], another equivalent construction combining two $\mathbf{C}\Phi(\delta)$ and single qubit rotations ($\alpha = x, y$)

$$e^{i(\delta/2)}\mathbf{ZZ}(\delta/2) = \begin{array}{c} \boxed{\mathbf{R}_x(\pi)} \text{---} \bullet \text{---} \boxed{\mathbf{R}_\alpha(\pi)} \text{---} \bullet \text{---} \\ | \quad | \\ \text{---} \boxed{\Phi(\delta)} \text{---} \boxed{\mathbf{R}_\alpha(\pi)} \text{---} \boxed{\Phi(\delta)} \text{---} \boxed{\mathbf{R}_x(\pi)} \text{---} \end{array} \quad (1.59)$$

is often employed to enable the range of negative and small angles.

Finally, on a 2-qubit quantum register it is straightforward to obtain $\mathbf{XX}(\delta)$ within \mathcal{S}_4 by directly using $\mathbf{T}_4(\delta, 0)$. Single qubit rotations can then be combined with this building block in a similar fashion as with other fundamental sets to obtain arbitrary 2-qubit evolution terms.

Unitary evolution terms of the form $\mathbf{U}_{\alpha\beta}^{(i,j)}(\delta)$ are readily generalized to N -qubit interactions

$$\mathbf{U}_{\alpha_1\dots\alpha_N}(\delta) = e^{-i\delta\otimes_i\sigma_{\alpha_i}^{(i)}} \quad (1.60)$$

In principle, such operations can always be decomposed into the single- and two-qubit gates, and this is usually the only solution available on processor architectures featuring only pairwise qubit interactions. An example within the \mathcal{S}_1 universal set is the following:

$$e^{-i\delta\sigma_z\otimes\sigma_z\otimes\sigma_z} = \begin{array}{c} \bullet \text{---} \bullet \text{---} \\ | \quad | \\ \oplus \text{---} \bullet \text{---} \bullet \text{---} \oplus \\ | \quad | \\ \oplus \text{---} \boxed{\mathbf{R}_z(2\delta)} \text{---} \oplus \end{array} \quad (1.61)$$

The pattern is naturally generalized to any $N > 3$, and changes of reference frames can be applied to individual qubits as done above for the $N = 2$ case [119]. In trapped ions processors, where the universal set \mathcal{S}_4 natively contains many-body interactions, the decomposition of N -body terms can usually be done very efficiently using Mølmer-Sørensen collective gates [115, 120] and the limits on N are in principle dictated only by the scalability of the hardware set-up itself.

1.1.4 Suzuki-Trotter approximation and digital error

The crucial ingredient enabling universal approximation of arbitrary quantum dynamics on a general-purpose quantum computer is the Suzuki-Trotter formula, Eq. (1.7). When the convergence towards the exact solution of the simulation is not trivial (e.g. when local Hamiltonian terms do not commute), the amount of acceptable digital error must be carefully assessed when designing the corresponding quantum algorithm. This is critical for intermediate-scale non error-corrected quantum




Figure 1.2: Comparison of different digital approximations for a 2-qubit TIM-like simulation. The solid black line shows the fidelity with respect to the exact simulation when a fixed number of ST steps $n = 5$ is used, failing after a very short phase evolution. The dotted red line tests a linear increase of n with δ according to $n = \delta/2\epsilon$, which also becomes inaccurate for larger phases. Finally, the solid green line shows the case in which the digital error is kept fully under control by increasing quadratically the number of ST steps according to $n = \delta^2/2\epsilon$. In the plot, $\epsilon = 0.1$ and n goes up to $n \simeq 10^4$ when the scaling is quadratic with the phase.

processors, where the increase in the number of gates which comes with the increase in the number n of Trotter steps cannot proceed indefinitely without affecting the quality of the results. In practical NISQ cases, it is usually sufficient for the digital error to be just smaller than the hardware noise.

To get a better understanding of the origin and scaling properties of the so called digital error, let us take two operators O_1 and O_2 such that $[O_1, O_2] \neq 0$. By exploiting the well known matrix exponential expansion

$$e^M = \sum_{n=0}^{\infty} \frac{M^n}{n!} \quad (1.62)$$

we can readily see that [121]

$$e^{(O_1+O_2)\delta} = \mathbb{I} + (O_1 + O_2)\delta + \frac{1}{2}(O_1^2 + O_1O_2 + O_1O_2 + O_2^2)\delta^2 + \mathcal{O}(\delta^3) \quad (1.63)$$

On the other hand,

$$\begin{aligned} e^{O_1\delta}e^{O_2\delta} &= \left(\mathbb{I} + O_1\delta + \frac{1}{2}O_1^2\delta^2 + \mathcal{O}(\delta^3)\right) \left(\mathbb{I} + O_2\delta + \frac{1}{2}O_2^2\delta^2 + \mathcal{O}(\delta^3)\right) \\ &= \mathbb{I} + (O_1 + O_2)\delta + \frac{1}{2}(O_1^2 + O_1O_2 + O_1O_2 + O_2^2)\delta^2 + \frac{\delta^2}{2}[O_1, O_2] + \mathcal{O}(\delta^3) \\ &= e^{(O_1+O_2)\delta} + \frac{\delta^2}{2}[O_1, O_2] + \mathcal{O}(\delta^3) \end{aligned} \quad (1.64)$$

Up to second order³ we then have

$$e^{O_1\delta}e^{O_2\delta} = e^{(O_1+O_2)\delta + \frac{\delta^2}{2}[O_1, O_2] + \mathcal{O}(\delta^3)} \quad (1.65)$$

³We are using here the fact, which can easily be checked by Taylor series expansion, that second order corrections in the formula $e^{x(A+B)} = e^{xA}e^{xB} + \mathcal{O}(x^2)$ are equal to the second order corrections in $e^{xA}e^{xB} = e^{x(A+B) + \mathcal{O}(x^2)}$ [121]. Higher order terms are different.

and we can also write

$$\begin{aligned} \left(e^{O_1 \frac{\delta}{n}} e^{O_2 \frac{\delta}{n}} \right)^n &= \left(e^{(O_1+O_2)\frac{\delta}{n} + \frac{\delta^2}{2n^2}[O_1, O_2] + \mathcal{O}\left(\frac{\delta^3}{n^3}\right)} \right)^n \\ &= e^{(O_1+O_2)\delta + \frac{\delta^2}{2n}[O_1, O_2] + \mathcal{O}\left(\frac{\delta^3}{n^2}\right)} \end{aligned} \quad (1.66)$$

Hence we obtain, again up to second order corrections, the so called first-order Suzuki-Trotter formula, see also Eq. (1.9),

$$e^{(O_1+O_2)\delta} \simeq \left(e^{O_1 \frac{\delta}{n}} e^{O_2 \frac{\delta}{n}} \right)^n - \frac{\delta^2}{2n}[O_1, O_2] \quad (1.67)$$

The inaccuracy $\epsilon = \delta^2/2n$ corresponding to the digital error grows quadratically with the phase of the evolution and decreases with the inverse of the number of digital ST steps. Higher order formulas for exponential product expansion are also known, such as the second-order Suzuki-Trotter [121]

$$e^{(O_1+O_2)\delta} = \left(e^{O_2 \frac{\delta}{2n}} e^{O_1 \frac{\delta}{n}} e^{O_2 \frac{\delta}{2n}} \right)^n + \mathcal{O}\left(\frac{\delta^3}{n^2}\right) \quad (1.68)$$

which provides a better scaling of the digital error at the cost of an additional factor per iteration. In all cases, a ratio $r_\epsilon = \delta^p/n^q$ controls the digital error as a function of the target evolution phase and the number of Suzuki-Trotter steps. Two different strategies can therefore be envisioned. On one hand, one could aim at a fixed digital precision ϵ over the whole range of the dynamical simulation. This requires to increase the number of Trotter steps, and consequently the total length of the quantum circuit to be computed, keeping the ratio r_ϵ fixed. As an example, for the first-order formula in Eq. (1.67) we get

$$n_\epsilon(\delta) \propto \frac{\delta^2}{2\epsilon} \quad (1.69)$$

Notice that while the number of digital steps increases, the phase evolution $\delta_n = \delta/n$ required in each step decreases as $1/n$, thus keeping the overall computation time on the physical hardware linear in the total phase provided that each digital step can be implemented with a coherent operation of duration $t \propto \delta_n$ [37]. On the other hand, when the maximum length of quantum circuits that can be faithfully realized is limited, such as in the case of NISQ processors, it might be convenient to keep the length of the quantum circuit (i.e. the number of steps n) fixed. This produces a phase-dependent digital error, scaling e.g. with δ^2 in the first-order case. The uniform effect of hardware noise over the whole simulation comes at the cost of a limited range of phases (and therefore of physical times) in which the results of the simulation accurately reproduce those of the target model. Hybrid solutions are also possible, e.g. by selecting reasonable number of steps n in different intervals of phases δ , always with the primary goal of balancing the total error arising both from the hardware noise and software-level approximations. An example of this approach will be given in a later chapter. The two different strategies (fixed ϵ vs fixed n) can be compared for example by computing the quantum fidelity

$$f(\delta) = |\langle \psi_0 | \psi_n(\delta) \rangle| \quad (1.70)$$

of the digitally evolved state

$$|\psi_n(\delta)\rangle = \left(e^{O_1 \frac{\delta}{n}} e^{O_2 \frac{\delta}{n}} \right)^n |\psi_0\rangle \quad (1.71)$$

with respect to the exact evolution

$$|\psi_{\text{ex}}\rangle = e^{(O_1+O_2)\delta} |\psi_0\rangle \quad (1.72)$$

The results are shown in Fig. 1.2 for a simple 2-qubit case where we choose $|\psi_0\rangle = |00\rangle$ and

$$O_1 = -i(\sigma_x^{(1)} + \sigma_x^{(2)}), \quad O_2 = -i\sigma_z^{(1)}\sigma_z^{(2)} \quad (1.73)$$

corresponding to a TIM-like interaction.

1.1.5 Extracting physical observables

The final target of a quantum simulation is the physical information contained in the evolved state of the quantum register $|\psi(t)\rangle$. In the final step of any simulation run, quantum measurements are applied to retrieve such information. In all practical cases, the easiest and often the only implemented kind of observation in NISQ processors is a measurement in the computational basis, i.e. of the occupation probability of the eigenstates of σ_z for each qubit. However, this is usually sufficiently versatile to give access to meaningful physical quantities. In simple cases, the expectation value of a generic quantum observable \mathbb{O} on the quantum register $\langle\mathbb{O}(t)\rangle = \langle\psi(t)|\mathbb{O}|\psi(t)\rangle$ can be reconstructed by a readout procedure combining unitary operations U_{meas} and measurements in the computational basis. Indeed, every \mathbb{O} is characterized by a set of eigenstates $\{|o_i\rangle\}$ and of possible corresponding outcome values $\{o_i\}$. It goes without saying that such \mathbb{O} must be written, under an appropriate mapping, in a form which can be interpreted on a quantum register, namely as a combination of Pauli operators. The role of U_{meas} is then to map eigenstates of \mathbb{O} onto computational basis states. As an example, if $\mathbb{O} = \sigma_x$, the readout of $\langle\sigma_x(t)\rangle$ on a single qubit can be done by performing a Hadamard gate, which map $\sigma_x \mapsto \sigma_z$, followed by a standard measurement in the computational basis. More in general, joint qubit measurements are also possible as a way of characterizing the output quantum state e.g. via quantum tomography [106]. Notice that a full reconstruction of $|\psi(t)\rangle$, which is a 2^N -dimensional complex vector, is actually exponentially expensive in terms of the number of repetitions of the simulation protocol: as a consequence, most quantum simulations aim directly at some specific physical property of the system under study.

A particularly useful strategy allowing the extraction of complex physical quantities and optimizing the efficiency of the measurement process is the ancilla-assisted method originally introduced by Somma *et al.* [40]. In view of the original applications that we are going to present in Chapter 2, we will now describe the ancilla-assisted observation of dynamical correlation functions and we will briefly comment how this can in turn lead to an algorithmic procedure for finding the spectrum of an Hermitian operator. Given a N -qubit state $|\psi\rangle$, a Hamiltonian H generating time evolution and two unitary operators V and W , we define the dynamical correlation $\mathcal{C}_{VW}(t)$ function as the quantity

$$\mathcal{C}_{VW}(t) = \langle V^\dagger(t)W \rangle = \langle \psi | e^{iHt} V^\dagger e^{-iHt} W | \psi \rangle \quad (1.74)$$

Dynamical correlations can be used to describe the propagation of excitations or external perturbations within a physical system. They appear in areas such as linear response theory [122], where they can be used to compute susceptibilities, magnetization and conductivity, manybody physics [123] and cross section calculations [124].

The ancilla-assisted quantum simulation of the dynamical correlation function can be performed with a quantum circuit whose basic structure is shown in Fig. 1.3. For simplicity, we assume that the quantum register is already prepared in the desired state $|\psi\rangle$ (it can be for example the ground state of the target physical system) and that the ancilla a starts in the quantum superposition $\frac{1}{\sqrt{2}}(|+\rangle) = |0\rangle + |1\rangle$. The joint initial state of the quantum register R and the ancilla is therefore $|\phi\rangle_{aR} = |+\rangle_a |\psi\rangle_R$. First, a W unitary is performed on R , controlled by the ancilla:

$$|\phi\rangle_{aR} \rightarrow \frac{1}{\sqrt{2}} (|0\rangle_a |\psi\rangle_R + |1\rangle_a W|\psi\rangle_R) \quad (1.75)$$

The sequence of gates for the digital simulation of the time evolution $U(t) = e^{iHt}$ is then applied to the quantum register, thus leading to

$$\frac{1}{\sqrt{2}} (|0\rangle_a U(t)|\psi\rangle_R + |1\rangle_a U(t)W|\psi\rangle_R) \quad (1.76)$$

Finally, a V unitary is applied to R , this time controlled by the state $|0\rangle$ of the ancilla⁴. The output state is:

$$|\phi_{out}\rangle = \frac{1}{\sqrt{2}} (|0\rangle_a VU(t)|\psi\rangle_R + |1\rangle_a U(t)W|\psi\rangle_R) \quad (1.77)$$

⁴By convention, controlled quantum gates are usually activated when the control qubit is in state $|1\rangle$. Activation conditioned by state $|0\rangle$ is easily obtained by adding $X \equiv \sigma_x$ quantum gates on a before and after the standard controlled operation.




Figure 1.3: Ancilla-based algorithm to compute dynamical correlation functions. The two alternative paths at the end of the circuit show a possible choice of unitaries U_{meas} which, followed by a measurement in the computational basis, give access to the real and imaginary parts of $\mathcal{C}_{VW}(t)$, proportional to $\langle \sigma_x \rangle$ and $\langle \sigma_y \rangle$ respectively.

A measure of the observable σ_x on the ancilla gives

$$\begin{aligned} \langle \sigma_x^{(a)} \rangle &= \text{Tr} \left[\left(\sigma_x^{(a)} \otimes \mathbb{I} \right) |\phi_{out}\rangle \langle \phi_{out}| \right] \\ &= \text{Re} [\mathcal{C}_{VW}(t)] \end{aligned} \quad (1.78)$$

In a similar way, $\text{Im} [\mathcal{C}_{VW}(t)]$ can be obtained by measuring $\langle \sigma_y^{(a)} \rangle$, in a second run of the algorithm. In total

$$\langle 2\sigma_+^{(a)} \rangle = \mathcal{C}_{VW}(t) \quad (1.79)$$

By removing the unitary evolution or by moving it at the beginning of the circuit, the same scheme can be used to simulate equal-time correlations. Notice that at the end of the proposed procedure all the information is accessible through the ancilla a alone, while the larger quantum register R needs not to be measured at all. This is, on a general basis, a good example of how physical information can be retrieved from a digital simulation without the need for full state tomography. Moreover, from a NISQ perspective this also means that the quality of the outcomes can be improved by mitigating, with suitable optimization techniques and processor design, the readout error of a single qubit. An evolution of the same algorithm can be used to efficiently extract n -point time-correlation functions [125] and the expectation values of any operator expressible in the form $\mathbb{O} = \sum_j c_j V_j^\dagger W_j$ where V_j, W_j are unitary operators [40].

Combining a slightly modified version of the ancilla-assisted strategy described above with either a classical Fast Fourier Transform (FFT) or even a Quantum Phase Estimation algorithm [5, 126] the spectrum of a Hermitian operator Q can be extracted. In particular, in physical problems one is usually interested in the case $Q = H$, where H is some Hamiltonian of interest. For example, the hybrid quantum-classical approach, proposed again by Somma *et al.* [40], requires the quantum register R to be initialized in a state $|\psi\rangle$ with some overlap with the eigenstates $|Q_l\rangle$ of Q

$$|\psi\rangle = \sum_l \lambda_l |Q_l\rangle \quad (1.80)$$

The exponential $U_Q(\theta) = e^{-iQ\theta}$ is a unitary operator which can be realized on the quantum register in exactly the same way as any standard time-evolution operator $U(\theta) = e^{-iH\theta}$. By setting e.g. $|\psi\rangle_R = |\psi\rangle$, $t = 0$ (i.e. removing the $U(t)$ part), $W = U_Q(\theta)$ and $V = \mathbb{I}$ in the algorithm in Fig. 1.3, we can then compute the expectation value $\langle \psi | U_Q(\theta) | \psi \rangle$. The result will be of the form

$$\langle U_Q(\theta) \rangle = \sum_l |\lambda_l|^2 e^{-i q_l \theta} \quad (1.81)$$




Figure 1.4: Quantum circuits for the 2-qubit Heisenberg model. (a) 6-CNOT decomposition within the \mathcal{S}_1 universal set. (b) Equivalent 3-CNOT decomposition [124, 127]. (c) 3- U_{XY} decomposition within the \mathcal{S}_2 universal set [53, 128].

where q_l are the eigenvalues of Q . By repeating the same procedure a number of times, the values of $\langle U_Q(\theta) \rangle$ can be reconstructed as a function of θ . The FFT then yields

$$\text{FFT}(\langle U_Q(\theta) \rangle) = \sum_l 2\pi |\lambda_l|^2 \delta(q - q_l) \quad (1.82)$$

An original combination of the quantum simulation of dynamical correlation functions with classical Fourier analysis, designed along the lines of what was presented in this section and run on a NISQ processor, will be presented and discussed in Chapter 2.

1.1.6 Examples

In order to see in action some of the techniques and procedures presented in the previous sections, we will now provide a few numerical examples applied to paradigmatic situations and models. The noise free classically computed results that we are about to show will constitute a sort of golden standard to assess the performances of real NISQ devices later on.

Recalling Eq. (1.22), we can write a 2-qubit isotropic Heisenberg model as

$$H_{\text{Heis},2} = J \left(\sigma_x^{(1)} \sigma_x^{(2)} + \sigma_y^{(1)} \sigma_y^{(2)} + \sigma_z^{(1)} \sigma_z^{(2)} \right) \quad (1.83)$$

The induced time evolution then reads

$$\begin{aligned} U_{\text{Heis},2}(\delta) &= e^{-i\delta(\sigma_x^{(1)}\sigma_x^{(2)} + \sigma_y^{(1)}\sigma_y^{(2)} + \sigma_z^{(1)}\sigma_z^{(2)})} \\ &= e^{-i\delta\sigma_x^{(1)}\sigma_x^{(2)}} e^{-i\delta\sigma_y^{(1)}\sigma_y^{(2)}} e^{-i\delta\sigma_z^{(1)}\sigma_z^{(2)}} \\ &= \text{XX}(\delta)\text{YY}(\delta)\text{ZZ}(\delta) \end{aligned} \quad (1.84)$$

where $\delta = Jt$ and the second equality, which is essentially the ST formula for $n = 1$, follows from $[\sigma_\alpha^{(1)}\sigma_\alpha^{(2)}, \sigma_\beta^{(1)}\sigma_\beta^{(2)}] = 0 \forall \alpha, \beta$. After identifying the two spin states $\{|\uparrow\rangle, |\downarrow\rangle\}$ with the two computational basis states $\{|0\rangle, |1\rangle\}$, a straightforward application of Eq. (1.48) and its generalization with single qubit rotations leads, within the universal set \mathcal{S}_1 , to a 6-CNOT decomposition of the Heisenberg time evolution, as shown in Fig. 1.4a. Despite its conceptual transparency, this is not the most effective solution to the 2-qubit Heisenberg model quantum simulation. Indeed, according to a result by Vidal and Dawson [127], any two qubit operation, therefore including $U_{\text{Heis},2}(\delta)$, can always be obtained with *at most* three CNOT gates and single qubit rotations. Based on such construction, a shorter but equivalent quantum circuit is reported in Fig. 1.4b. Similarly, in \mathcal{S}_2 we can obtain the digital simulation of $U_{\text{Heis},2}(\delta)$ either by juxtaposing gate sequences of the form shown in Eq. (1.55) or by using only three 2-qubit U_{XY} gates (instead of six), as reported in Fig. 1.4c. This last result is based on the identity [53, 128]

$$H_{\text{Heis},2} = \frac{J}{2} (H_{xyxy} + H_{xxzz} + H_{zzyy}) \quad (1.85)$$




Figure 1.5: Digital quantum simulation of spin models. Exact curves are obtained with full Hamiltonian exponentiation, the dotted lines represent ideal digital approximation computed with the Suzuki-Trotter formula, and the quantum simulations data points, indicated with dots and diamonds, are computed multiplying the matrix representation of the corresponding sequences of elementary quantum gates. Increasing degrees of digital resolution, involving longer underlying quantum gate sequences, provide in principle better agreement with the ideal solution of the target model. (a) Individual spin magnetization for the 2-qubit Heisenberg model, using the decomposition in Fig. 1.4b for the digital quantum simulation. The initial state of the two spins is $\sqrt{2}|\psi_0\rangle = |\uparrow\rangle(|\uparrow\rangle + |\downarrow\rangle)$. (b) Time evolution of the occupation probability of the initial state $|\psi_0\rangle = |100\rangle$ of 3 qubits interacting as a Heisenberg chain with open ends in an external field. Here $J_{12} = J_{23} = J$ and $Bg = 20J$. (c) Total magnetization of a pair of qubits interacting according to the Transverse Field Ising model (TIM), with $J_{zz} = J$ and $Bg = 2J$. The digital quantum simulation is performed using the matrix representation of the \mathcal{S}_1 fundamental set.

where $H_{\alpha\alpha\beta\beta} = \sigma_\alpha^{(1)}\sigma_\alpha^{(2)} + \sigma_\beta^{(1)}\sigma_\beta^{(2)}$ and all the terms on the right hand side commute. Just from these two elementary examples, we have once more a demonstration of the fact that no absolute solution is intrinsically valid for the translation of a certain physical time evolution in a sequence of digital quantum gates, not even within the same fundamental set of operations. Optimization can be performed either in the design of the algorithm or, on a more technical level, at the compilation stage which adapts a certain quantum circuit to the physical properties of a real quantum processors. On NISQ devices, it is often crucial to have the shortest possible circuits and the smallest number of two-qubit operations: however, several other constraints might enter the game, such as qubit quality or chip connectivity. In \mathcal{S}_3 a decomposition with three $C\Phi(\delta)$ follows immediately from Eq. (1.58) and single qubit changes of reference frame. Finally, in \mathcal{S}_4 a possible realization of the Heisenberg interaction can be obtained for some digital resolution δ as

$$U_{\text{Heis},2}(\delta) = \text{ABCAC}^\dagger \quad (1.86)$$

where $A = T_4(\delta, 0)$, $B = T_4(\delta, \pi/2)$ and $C = T_3(\pi/4, \pi/2)$. With any of the above elementary decompositions, the digital quantum simulation of the 2-qubits Heisenberg model can be performed and physical information can be extracted by using the methods discussed in Sec. 1.1.5. A numerical example computed in Matlab is reported in Fig. 1.5a, where we show the digital quantum simulation of the individual magnetization of the two spin-1/2 particles. The latter can be extracted by measuring the observable $\sigma_z^{(i)}$ and using the definition $\langle s_z^{(i)} \rangle = (1/2)\langle \sigma_z^{(i)} \rangle$. No digital error is present in this case, due to the commutativity of the local terms in the 2-qubit Heisenberg chain.

Simple 2-spin chains such as the Heisenberg bond just described can be used as building block to realize more complex configurations, thus extending the quantum simulation to arbitrary number of spins with pairwise interactions and different inter-qubits connectivity. For example, a 3-spin Heisenberg chain with open ends, put in an external magnetic field \vec{B} , has a Hamiltonian of the form

$$H_{\text{Heis},3} = H_B + H_{\text{Heis},2}^{12} + H_{\text{Heis},2}^{23} \quad (1.87)$$

where

$$\mathbf{H}_B = \frac{Bg}{2} \left(\sigma_z^{(1)} + \sigma_z^{(2)} + \sigma_z^{(3)} \right) \quad (1.88)$$

describes the magnetic field oriented along the z -direction and each of the spin-spin bonds corresponds to a term

$$\mathbf{H}_{\text{Heis},2}^{ij} = J_{ij} \left(\sigma_x^{(i)} \sigma_x^{(j)} + \sigma_y^{(i)} \sigma_y^{(j)} + \sigma_z^{(i)} \sigma_z^{(j)} \right) \quad (1.89)$$

In general, the two bonds can be nonequivalent, i.e. $J_{12} \neq J_{23}$. Since $[\mathbf{H}_{\text{Heis},2}^{12}, \mathbf{H}_{\text{Heis},2}^{23}] \neq 0$ (independently from the coupling constants J_{ij}), in this case the quantum simulation must be carried out using the ST digital procedure, alternating the application of the sequence for the 2-spin case on the two bonds

$$\mathbf{U}_{\text{Heis},3}(\delta) = \left(\mathbf{U}_{\text{Heis},2}^{12}(\delta_{12}/n) \mathbf{U}_{\text{Heis},2}^{23}(\delta_{23}/n) \right)^n e^{-i\mathbf{H}_B t} \quad (1.90)$$

where $\delta_{ij} = J_{ij}t$. The part describing the magnetic field on equivalent spins (we set the gyromagnetic ratio $g_1 = g_2 = g_3 = g$) corresponds to single qubit rotations around the z axis. Since this part commutes with the rest, it can be performed at the beginning of the circuit and does not enter in the Suzuki-Trotter discretization procedure. In Fig. 1.5b we show how these results can be used to compute the time evolution of the occupation probability of an initial state $|\psi_0\rangle = |\uparrow\downarrow\downarrow\rangle = |100\rangle$. This approach can be now straightforwardly generalized to more complicated spin lattices. Remarkably, all bond interaction involving disjoint subsets of spins generate locally independent time evolution terms commuting with each other: indeed, by construction of the tensor product space, $[\sigma_\alpha^{(i)}, \sigma_\beta^{(j)}] = 0$ when $i \neq j$. As a consequence, the corresponding unitary operators can always be simulated in parallel on the quantum register and decomposed in just as many elementary quantum gates as in simpler configurations, thus reducing the overall complexity of the quantum simulation. As we extensively discussed, this mechanism lies at the very heart of the computational efficiency of digital quantum simulation algorithms.

Other paradigmatic forms for pairwise spin interactions are also very easily obtained with similar techniques. For example, we recall that the Hamiltonian of the Transverse Field Ising model (TIM), introduced in Eq. (1.24), in the two-qubit case can be written as

$$\mathbf{H}_{\text{TIM},2} = \mathbf{H}_{B,x} + \mathbf{H}_{zz} \quad (1.91)$$

where

$$\mathbf{H}_{B,x} = \frac{Bg}{2} \left(\sigma_x^{(1)} + \sigma_x^{(2)} \right) \quad \mathbf{H}_{zz} = J_{zz} \sigma_z^{(1)} \sigma_z^{(2)} \quad (1.92)$$

The quantum simulation of the TIM corresponds to the following digital process

$$\mathbf{U}_{\text{TIM},2}(t) = \left(\mathbf{ZZ}(J_{zz}t/n) e^{-i(\sigma_x^{(1)} + \sigma_x^{(2)}) \frac{Bgt}{2n}} \right)^n \quad (1.93)$$

Apart from straightforward single qubit rotations around the x axis, the required quantum circuit contains only ZZ operations, which can easily be translated into elementary quantum gates as done, e.g., in Eq. (1.48). In Fig. 1.5c we report the time evolution of the total magnetization of the spin dimer along z , simulated with the matrix representation of the quantum gates in \mathcal{S}_1 and extracted by measuring the expectation values of $\sigma_z^{(i)}$.

Finally, we also report an example featuring the ancilla-based algorithm discussed in Sec. 1.1.5. Here, it is applied to extract spin-spin dynamical correlations

$$\mathcal{C}_{ij}^{\alpha\beta}(t) = \langle s_\alpha^{(i)}(t) s_\beta^{(j)} \rangle = (1/4) \langle \sigma_\alpha^{(i)}(t) \sigma_\beta^{(j)} \rangle \quad (1.94)$$

for the 3-spin open Heisenberg chain of Eq. (1.87). Dynamical correlations are computed on the ground state of the system, which for a sufficiently strong external magnetic field is well approximated by $|\psi\rangle = |\downarrow\downarrow\downarrow\rangle$. The structure of the required quantum circuit is shown in Fig. 1.6a for the case of next-to-nearest neighbors cross correlations. Autocorrelations and nearest neighbors correlations can be computed in a similar way by changing the target qubit involved in the operations controlled by the ancilla. Numerical results based on \mathcal{S}_1 decompositions are presented in Fig. 1.6b-d.




Figure 1.6: Digital quantum simulation of dynamical correlation functions for the 3-qubit Heisenberg model. The solid line is the exact computation without digital errors, the dotted line represents the expected result for continuous phase Trotter evolution with $n = 5$ and the data points show the result of the corresponding quantum circuit for a selection of phase values. Here we set $J_{12} = J_{23} = J$ and $Bg = 20J$. (a) Quantum circuit allowing to compute time-correlation functions between next-to-nearest neighbors qubits, namely $\langle \sigma_\beta^{(3)}(t) \sigma_\alpha^{(1)} \rangle$. The operators α and β represent (controlled) σ_α and σ_β unitary transformations, $\delta_B = Bgt$ and $U_{\delta/n}^{ij}$ is a shorthand notation for $U_{\text{Heis},2}^{ij}(\delta_{ij}/n)$. The part inside the green box must be repeated n times. (b) Autocorrelation $\langle s_x^{(1)}(t) s_x^{(1)} \rangle$. (c) Nearest neighbors $\langle s_x^{(2)}(t) s_x^{(1)} \rangle$ cross correlation. (d) Next-to-nearest neighbors $\langle s_x^{(3)}(t) s_x^{(1)} \rangle$ cross correlation.

1.2 State-of-the-art in digital quantum simulations

Having laid the theoretical foundations of digital quantum simulations, and before presenting in depth some original results, we will devote the final part of this chapter to a brief overview of the main experimental achievements constituting the state-of-the-art at the time of writing. While it is certainly difficult to keep up with the most recent studies in years of excitement and fast progress, we will nevertheless mention a few seminal papers who are unanimously recognized as the ones who established the field as a whole. This section is mainly focused on the two different technological approaches which have so far been effectively demonstrated as promising NISQ processor candidates, i.e. trapped ions and superconducting circuits. Even though a thorough comparison between the two is inevitably affected by very specific experimental details, we will try to give a broad and synthetic overview of the respective strengths and weaknesses. An outlook on possible future platforms for universal quantum simulations and quantum information processing is left for a separate discussion in Chapter 3.

1.2.1 Trapped Ions

Since the late nienties, and in parallel with the impressive advancement in the field of manipulation and control of single quantum systems [129, 130], atomic ions in linear Paul traps [131] have been representing one of the most promising candidates for realizing fully operational quantum processors [68, 74, 132, 133]. As compared to neutral atoms, the trapping potential for charged atomic species can be much stronger, thus allowing to hold each single ion for several hours, and even days, with very long coherence times and a remarkably high degree of external control. This makes them a reliable solution for quantum information processing satisfying all of the required DiVincenzo criteria. As of today, digital quantum processors with up to 11 programmable ions in a linear Paul trap have been made available [134], with some spin-off companies such as IonQ, Honeywell and Alpine Quantum Technologies offering restricted cloud access⁵.

Despite the technical challenges posed by ultra-high vacuum and laser cooling, trapped ions processors do not in general need cryogenic temperature to be operated, which represents a significant advantage in terms of costs and portability over competing technologies. In this hardware platform, individual ions are held in place through radio frequency oscillating electric fields that generate a stable two-dimensional potential well [135]. Ions are then confined by a further harmonic trap along a line and, due to the mutual Coulomb repulsion, they almost spontaneously arrange in a chain whose nodes are spatially separated by a few microns. While this is currently the most performing configuration, it is nevertheless considered hardly scalable towards larger sizes or higher dimensionality of the confining potential, as the number of ions in a chain affects their separation and, ultimately, limits the capabilities of individual qubit control and the multi-qubit gate fidelity. In this respect, modularity [136] and 2D arrays [137] are being considered as possible solutions.

Different types of trapped ion qubits can be defined, depending on the frequency spacing between the relevant energy eigenstates selected to encode the logical $\{|0\rangle, |1\rangle\}$ basis. In general, hyperfine qubits are encoded into a pair of energy levels typically separated by frequencies in the GHz range, while optical qubits are defined corresponding to quadrupole active transitions in the hundreds of THz. Zeeman qubits can also be defined via the application of an external static magnetic field, opening a low frequency (few MHz) tunable gap between magnetic levels. The technologies based on $^{40}\text{Ca}^+$ as optical qubits [138, 139], and $^{43}\text{Ca}^+$ or $^{171}\text{Yb}^+$ as hyperfine qubits [140, 141], respectively, are particularly advanced, although several other atomic species with a single outer electron can be successfully trapped [74, 133].

Irrespective of the specific qubit realization, initialization, control and readout are performed via coherent manipulation. For example, initial state preparation can be obtained using optical pumping with external lasers of suitable frequency, while readout is achieved by detecting resonantly scattered radiation from an optical transition. Single-qubit operations are enabled by directly coupling the designated $|0\rangle$ and $|1\rangle$ eigenenergy levels, and the required tools inevitably depend on the qubits type [135]. Multi-qubit gates between ions trapped along the same chain are realized by exploiting the transverse normal vibrational modes of the whole ion string trapped in a harmonic potential [142], which are used as a bus to transfer quantum information. One of the most common native entangling logical port relies on a controlled-phase type of gate that was originally proposed from Mølmer and Sørensen [116], see Eq. (1.42). The details of the specific implementation again depend on the particular type of qubit [75, 117, 118]. We notice that an unequivocal advantage of this set-up, which is often of crucial importance for quantum simulation purposes, is that the design of the processor features a built-in all-to-all connectivity between qubits: no other alternative technology has so far reached comparable levels of flexibility with respect to chip topology.

In terms of absolute performance, the different experimental platforms have essentially shown quite comparable figures of merit. In particular, when working with isolated or pairs of trapped ions, single-qubit operations with fidelities in the order of 99.9999% have been reported [143], as well as two-qubit gates reaching fidelities above 99.9% even in different experimental setups [140, 144]. Typical duration for single-qubit gates varies between 100 ns and few tens of μs , while two qubit gating times range in the μs to few hundred μs interval [133, 140, 145]; readout is typically performed in hundreds of μs with fidelities in the 99.99% range [146].

⁵See e.g. <https://ionq.com/>, <https://www.honeywell.com/en-us/company/quantum> and <https://www.aqt.eu/>.

Since the typical coherence times of trapped ion qubits vary between few hundred of milliseconds [147] to hundreds of seconds [143, 148], the trapped ion quantum hardware is the one currently allowing to achieve the highest coherence versus gating time ratio, ideally in the order of 10^5 to 10^6 . In view of practical quantum computing applications, we should however emphasize that the considerations above are mostly limited to few qubits quantum hardware: while analog quantum simulators of Ising chains have been shown with more than 50 trapped ions [35], the stringent performance requirements for digital quantum information processing are still only reachable with NISQ prototypes with no more than 20 qubits [138]. In fact, increasing the number of ions in the chain ultimately limits gates fidelity, mainly due to the difficulties of individually addressing single qubits while avoiding cross-talks between the different beams. Moreover, errors affecting the overall success of a quantum computation generally arise from two distinct mechanisms: decoherence, induced e.g. from undesired qubit-environment coupling such as spontaneous emission, frequency shift, motional heating or spurious residual collisions inside the vacuum chamber, and imperfect control fields, such as miscalibrated or noisy control field amplitude, frequency, or polarization, which typically result in quantum gate errors [133]. On a quantitative level, it has recently been reported that the fidelity of multi-qubit entangling Mølmer-Sørensen gates degrades from 99.6% with 2 optical qubits down to 86% when the same quantum hardware is loaded with 10 optical qubits [149]. At the same time, a quantum processor based on 11 hyperfine qubits has recently been shown to achieve an all-to-all connectivity with average single- and two-qubit XX gate fidelities of 99.5% and 97.5% respectively [134].

Probably the first experimental demonstration of digital quantum simulations came in 2011, performed by Lanyon *et al.* [115] on a trapped ions prototype quantum processor. Some of those early results on paradigmatic few spin toy-models still stand as reference standards. As an example, we report in Fig. 1.7 the digital simulation of two-spin Ising, XY and XYZ models: the most striking feature, which had never been proved before, is of course the capability of the same quantum hardware to be reprogrammed to simulate essentially different interaction terms with no direct connection to the physical Hamiltonian of the setup. Each spin-1/2 was directly mapped onto a single ion, and the unitary operations reported in Fig. 1.7

$$C = T_2(\pi/16) \quad D = T_4(\pi/16, 0) \quad E = T_4(\pi/16, \pi) \quad F = T_3(\pi/4, 0) \quad (1.95)$$

were defined in terms of the universal \mathcal{S}_4 set of native gates, see Eqs. (1.40)-(1.42). The time evolution is quantified by a dimensionless phase $\theta = Et/\hbar$ and was obtained through several repetitions of Suzuki-Trotter digital steps at fixed resolution⁶. The initial state was chosen as an eigenstate of $\sum_i \sigma_x^{(i)}$, and the population in each of the eigenstates was monitored as a function of θ . The same work reported remarkable results for the digital quantum simulation of up to 6 spins and multi-spin interaction terms, thus showing from the beginning the superior potential of trapped ions architectures in approaching such complex models.

Going beyond spin models, the digital quantum simulation of lattice gauge field theories was implemented by Martinez *et al.* [61], in a neat demonstration of the possible broad impact quantum simulation techniques can have on different areas of physics. In this work, Kogut–Susskind fermionic degrees of freedom were encoded into Pauli operators, while the electromagnetic background was eliminated via the gauge freedom of the theory. The authors reported a successful simulation of particle-antiparticle spontaneous creation from vacuum fluctuations (i.e. the Schwinger mechanism) and of the corresponding entanglement dynamics on a 4-qubit proof-of-principle realization. Open

⁶Notice that this is yet another possible strategy to tackle the issue of the digital error, see Sec. 1.1.4. Here, a fixed size of the phase slice $\delta_n = \delta/n$ is chosen, essentially following the idea of $n \propto \delta$ as in one of the cases presented in Fig. 1.2. For many practical purposes, the exact dynamics is well approximated with this approach: indeed, with the same steps leading to Eq. (1.67) we can write

$$e^{(O_1+O_2)\delta} \simeq \left(e^{O_1\delta_n} e^{O_2\delta_n} \right)^{\delta/\delta_n} + \mathcal{O}(\delta_n\delta)$$

The agreement with the exact dynamics is of course better for smaller δ_n and is controlled by the product $\delta_n\delta$. From an experimental perspective, having a fixed phase evolution makes it possible to use the same pulse scheme for each step, such that only one or few control sequences must be optimized.




Figure 1.7: Quantum simulations with a trapped ions processor. Experimental quantum simulation of two-spin models: the complexity of the simulated model increases from A to C, the digital resolution is fixed to $\theta/n = \pi/16$ and the number of digital steps increases linearly with the simulated phase evolution. In each panel, a sketch of the sequence of unitary operators corresponding to a single digital step is shown, and up to $n = 12$ consecutive Suzuki-Trotter steps were performed for the data shown here. Solid lines correspond to the exact evolution, empty symbols correspond to the ideal digitized evolution, and filled symbols are the quantum simulator results for the evolution of the different eigenstates. Reproduced from Ref. [115].

quantum systems dynamics and quantum maps have also been approached on trapped ions processors [114, 120].

With respect to quantum register size, a fully controlled 20-qubit system has been shown to reliably allow for the creation of multi-qubits entangled states [138], and a remarkable hybrid quantum-classical simulation of the Schwinger model on 20 $^{40}\text{Ca}^+$ ions was very recently reported by Kokail *et al.* [150]. Here, digitalized Hilbert space exploration techniques such as the Variational Quantum Eigensolver (VQE) algorithm [151] were used to study static ground state properties and phase transitions. The VQE algorithm has also been applied to study from very simple [139] to more complex chemical species, such as the water molecule [152], using in the latter case a hyperfine qubits-based quantum processor. Along the same lines, the deuteron nucleus binding energy within a percent accuracy was also obtained with a ion trap quantum processor implementing a record quantum circuit depth [153], thus paving the way to possible further studies of effective field theories in nuclear physics.

1.2.2 Superconducting circuits

From the early days of circuit Quantum Electrodynamics (cQED) [76, 77, 154, 155], a fast and continuous technological improvement [60, 154–162] has brought superconducting quantum circuits to challenge trapped ions as the leading platform for practical quantum information processing [98]. Given their inherent conceptual simplicity and relatively close relationship with many standard electrical engineering techniques, it is not surprising that even outside academic laboratories some of the worldwide leading and emerging high-tech companies currently investing in quantum computing are concentrating their efforts on this architecture (e.g. Google, Rigetti and IBM, the latter with the IBM Q program).

Several thorough reviews about the technical implementation of superconducting qubits, describing different design solutions and performances, can be found in the literature [72, 78, 109]. Superconducting circuits require cryogenic temperatures around 10-15 mK to be reliably operated in the quantum regime, a condition that is often achieved in $^3\text{He}/^4\text{He}$ dilution refrigerators. Qubits are encoded into the energy spectrum of the lowest collective charge/current excitations in a micro-LC resonator, with a nanostructured Josephson junction playing the role of a nonlinear inducting element. By changing the relative values of characteristic circuital parameters, such the charging energy of a single Cooper pair, the Josephson or the inductive energy, different families of superconducting qubits can be identified, such as phase qubits, rf-sQUID, flux qubits, charge qubits and transmons. In particular, the latter are essentially a capacitively shunted version of the basic anharmonic LC resonator devices [156] and can reach coherence times in the 100 μs range [161], with an increase of up to 5 orders of magnitude with respect to the original Cooper Pair Box design from the late Nienties. These impressive results have even called for a cQED analog of Moore’s law, named after R. J. Schoelkopf, stating that the best coherence times available on superconducting qubits increases by an order of magnitude every three years [78]. In view of its very good properties, the transmon is nowadays the elementary unit of several scalable proposal for quantum computing architectures [86], and other transmon-inspired qubits, such as the Xmon [60] have also shown remarkable fidelities in a setup consisting of 9 qubits [163, 164]. A transmon processor consisting of 53 operating qubits has been recently reported [83] by the Google AI Quantum research group, showing improved single- and two-qubit gate fidelities.

Transmission lines wired at the edges of the chip board allow for individual qubit addressability, thus enabling single-qubit initialization, manipulation and read-out through microwave control pulses. In particular, current superconducting quantum circuits allow for single-qubit gate fidelities above 99.9% [60]. Qubit-qubit couplings are obtained through additional superconducting transmission line resonators, and two major alternative schemes are employed to engineer two-qubit quantum logic gates: in one case, the qubit transition frequencies are tuned by a local magnetic field to activate the mutual interactions, while in the second version a cross-resonant (CR) drive is used, applying to a given control qubit microwave pulses in resonance with another target qubit connected to the previous one. While the first realizations of the frequency tuning scheme were obtained by bringing the two qubits into resonance to get a virtual photon exchange, yielding an effective XY interaction [106], the most promising implementations are currently based on tuning one qubit along a “fast adiabatic trajectory” [60] that moves the $|11\rangle$ component of the wave-function close to an avoided level crossing with state $|02\rangle$, leading to a state-dependent phase and hence to the implementation of the controlled-phase gate [165]. This approach results in very fast (typical gate duration around 40 ns) and high-fidelity two-qubit gates, compared to relaxation and coherence times in the 20-40 μs range on average [60]. Conversely, the CR scheme uses fixed-frequency qubits in order to avoid frequency crowding and to reach longer coherence times up to 50 μs ⁷ and can easily implement a CNOT gate [104]. However, as the CR effect is essentially perturbative (see also Sec. 2.1), this requires more selective pulses and hence results in slower two-qubit gates, taking $\sim 200 - 300$ ns on average [104, 166], with average fidelity of 96%. A novel proposal to directly implement exchange-type gates with tunable amplitude and phase on fixed frequency qubits, thus making this architecture much more flexible e.g. for quantum simulations applications, has also been recently reported [167].

⁷Since the frequencies are fixed, the qubits can for example be operated constantly at a coherence “sweet spot”.




Figure 1.8: Digital quantum simulations on a superconducting circuit processor. (Top) Experimental quantum simulation of the Ising model in a transverse homogeneous field for two spins with increasing number of Trotter steps. (Bottom right) Dependence of final state fidelity on the number of digital steps used in the quantum simulation, for different phase angles (color bars), compared to ideal unitary evolution for the given number of Trotter steps. (Bottom left) False color image of superconducting quantum processor with 4 Niobium qubits and Aluminum transmission line resonators. Output and output ports and single qubit flux bias lines are also highlighted. Reproduced from Ref. [106].

As already noted for trapped-ion architectures, it is worth pointing out that while remarkable results have been obtained in samples specifically aimed at testing the basic operations on few qubits circuits, obtaining a scalable platform able to implement fast and high fidelity single- and two-qubit gates, as well as efficient readout, still represents a key challenge. For example, we will extensively describe in Chapter 2 how systematic coherent inaccuracies due to an imperfect implementation of the elementary gates are nowadays one of the leading error sources when several qubits are operated together to practically realize digital quantum simulation protocols.

The use of superconducting quantum hardware for digital quantum simulations of spin models was first reported in 2015 by Salathé *et al.* [106]. In their work, the authors successfully demonstrated, on a 4-qubit quantum processor, proof-of-principle realizations of 2-spin Heisenberg and Ising models, with several digital steps performed in sequence. As already briefly mentioned when introducing the universal operation set \mathcal{S}_2 , see Eq. (1.34), superconducting processors with frequency-tunable qubits coupled through an intermediate resonator naturally implement a XY-type interacting spin Hamiltonian, which can be used as the basis to digitally program a full Heisenberg or Ising type evolution, see for example Fig. 1.4c and Refs. [53, 106]. In Fig. 1.8 we report typical results for the digital evolution of the spins projections along the magnetic field direction, z , with up to 3 digital time steps, for an initial state prepared in $|\uparrow\rangle$ ($(|\uparrow\rangle - i|\downarrow\rangle)/\sqrt{2}$) that evolves non trivially in time according to the Transverse Fied Ising Hamiltonian. In another panel of Fig. 1.8, a summary of the fidelity for the quantum simulation with up to 5 Trotter steps is also reported. As it can be observed, while the ideal fidelity of the simulated quantum state with respect to the exact evolution increases against the number of Trotter steps, the experimental one starts to decrease after about

2 or 3 digitized steps, depending on the phase, due to the increased length and complexity of the underlying quantum circuit. Despite the fact that approximately 5 ST steps still represented an upper bound beyond which the quality of the results is heavily affected by short coherence times and systematic gate errors, the results presented here have undoubtedly set a standard in the field of cQED-based quantum information processing for physics applications, and they still stand as a touchstone for digital quantum simulations in superconducting quantum circuits. Last but not least, in the same work a first attempt at simulating the digital time evolution of two-point correlations was also reported.

As with trapped ions processors, the potential usefulness of the superconducting circuit quantum hardware as a universal quantum simulator goes well beyond the analysis of bare spin models. For example, other seminal works [93, 168] reported the experimental digital simulation of 3 and 4 modes Fermi-Hubbard model, a notoriously difficult model to be addressed with classical algorithms. More recent results in hybrid quantum-classical (e.g. VQE) approaches have given a boost to the field of quantum chemistry [94, 169], and nuclear physics problems have also been tackled with superconducting quantum processors. In particular, an evolution of the VQE algorithm led to the cloud-based computing of the deuteron binding energy [95] on IBM Q public devices. Moreover, a quantum-classical algorithm has been used to approach the solution of the Schwinger model dynamics [62]. In order to partially overcome at least some of the barriers hindering the progress towards a first demonstration of quantum advantage on NISQ processors, considerable work is currently dedicated to understanding the main sources of error and the noise properties on superconducting processors. Error mitigation techniques, which are being discussed, designed and tested, have shown the potential for improving the overall quantum simulation fidelities [62, 124, 170–172].

1.2.3 Summary and comparison

In this section we have presented an inevitably partial but representative selection of results demonstrating the state of the field of practical digital quantum simulations on NISQ devices. In Fig. 1.9, we try to give a quantitative summary of the main achievements in terms of fidelity and number of digital steps. Such a direct comparison has to be taken with a grain of salt, as it is certainly difficult to put in the right perspective experiments and works performed on different platforms, under different initial conditions, and reporting slightly different figures of merit. Nevertheless, the plot gives a visual idea of the scenario on digital quantum simulation of spin Hamiltonians up to date. Despite the clear and still considerable correlation between the number of digital steps included in the simulation and the fidelity of the final state obtained, we can see that currently trapped ions quantum simulators allow better performance for deeper quantum circuits, i.e. with large numbers of Trotter steps. With respect to this, we also recall that up to 12 digital steps have been achieved on trapped ions processors [115], albeit such data point is not reported here due to the lack of the fidelity characterization in the original publication. We also notice that a number of 5 Suzuki-Trotter digital steps is currently a limiting value for superconducting quantum processors if the fidelity of the simulation is to be kept above a reasonable accuracy threshold, meaning that there is still much room for improvement. In terms of the size of the simulated model, pure digital quantum simulations for spin-models with up to 6 spins have been performed on trapped ions processors [115]. On the other hand, in Chapter 2 we will present original results on superconducting processors with up to 4 spins [124], which is currently close to the technological limits of the platform. All in all, the picture emerging from the data collected here is in line with recent studies comparing the two architectures when challenged with similar quantum algorithms on 5 qubits processors [98]. We should stress once more that most of the remarkably high experimental fidelities reported in the literature have been achieved on few-qubit setups that, although scalable by design, are optimized for a specific target and are often pushed to the experimental limits. Conversely, when several qubits are connected and operated together, many new challenges emerge, such as the need to selectively address only some of them or to keep cross-talks under control.

On the technical side, the two leading architectures have shown many complementary characteristics. While superconducting circuits usually achieve much larger gate speeds, and thus may offer better performances when comparing quantum and classical devices in terms of absolute execution

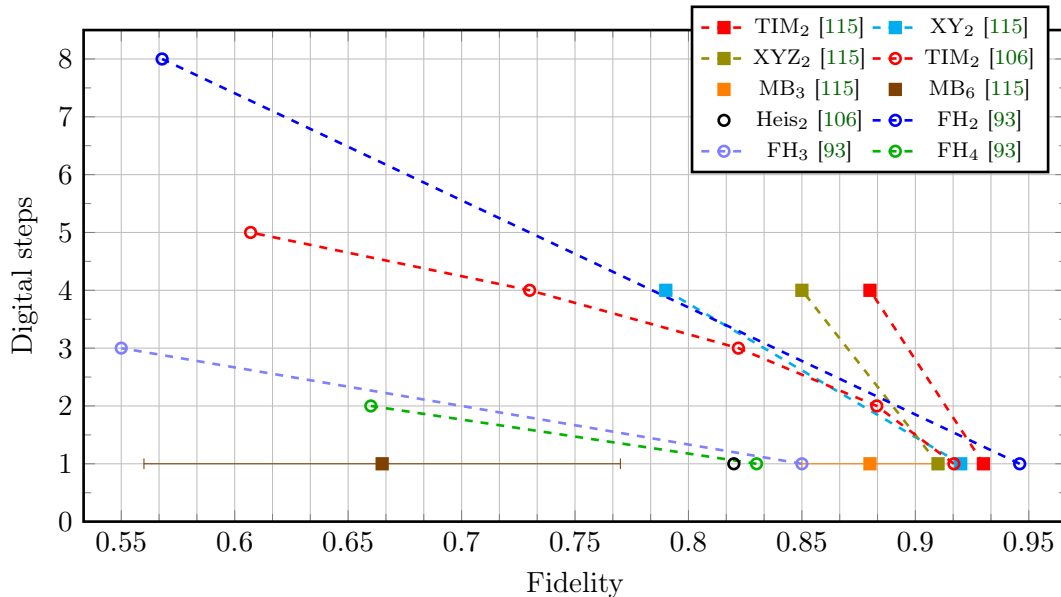


Figure 1.9: Summary of state-of-art experimental digital quantum simulations. Open circles represent results obtained on superconducting circuits quantum processors, while squares correspond to experimental quantum simulations on trapped ions processors. The color code corresponds to different target models being simulated: two-spin Transverse field Ising model (TIM₂), two-spin XY (XY₂) and XYZ (XYZ₂) models, 3- and 6-spin many body interactions (MB_{3,6}, fidelities given as estimated bounds), two-spin Heisenberg model (Heis₂), and 2- to 4-mode Fermi Hubbard model (FH_{*x*}, with $x = 2, 3, 4$). Although more digital steps than the ones reported here were actually performed in some of the experiments, data points are shown only when some measure of accuracy was provided in the original reference. Fidelities from Ref. [106] are given with respect to the ideal evolution for a fixed phase value and initial state, while those from Ref. [115] are process fidelities given with respect to the expected digitized evolution. Finally, data from Ref. [93] are extrapolated linear trends of fidelity with respect to the ideal digital outcome.

time of a given algorithm, the ratio of coherence time to gate operation remains in favor of trapped ions architectures. For digital quantum simulation purposes on NISQ devices, a native all-to-all connectivity is often a significant advantage, as it effectively reduces typical circuit depths and helps in avoiding the use of cumbersome SWAP gates: in this sense, trapped ions are more promising for establishing quantum correlations between distant pairs of qubits, although the first attempts to go beyond nearest-neighbors coupling have been reported on three-qubit superconducting quantum processors [173]. Finally, while in trapped ion-based technologies all qubits are in principle identical, superconducting qubits usually display, even on the same chip, different fabrication parameters and qualities. These are also affected by thermal cycling and hence require a detailed and frequent characterization to accurately calibrate the control pulses [174].

In the interest of finding platform- and problem-independent figures of merit to compare different quantum computing architectures, the so called *quantum volume* [99, 175, 176] has been introduced in recent years to provide combined information both on the available quabits N and the typical circuit depth $d(N)$ that can successfully be run on a given platform. Conceptually, the quantum volume is defined as

$$V_Q = \min[N, d(N)]^2 \quad (1.96)$$

To be more specific, let us denote by $\epsilon_{\text{eff}}(N)$ the average error rate per two-qubit general SU(4) gate run on any pair of qubits in a N qubit quantum register on the hardware under examination. Notice that $\epsilon_{\text{eff}}(N)$ in general receives contributions from hardware-native single- and two-qubit gate errors, from the possible lack of all-to-all connectivity and from the structure of the available fundamental gate set. As a result, $\epsilon_{\text{eff}}(N)$ is typically different, and usually larger, than the ideal error rate measured for proof-of-principle demonstrations of individual gates. We can then assume that the

error rate per clock cycle, i.e. per step in any given algorithm when all possible parallelization of two-qubit operations on disjoint pairs of qubits is taken into account, scales as $\epsilon \simeq N\epsilon_{\text{eff}}(N)$. An estimate of the average circuit depth in which at most a single error occurs now gives

$$d \cdot \epsilon \simeq 1 \implies d \simeq \frac{1}{N\epsilon_{\text{eff}}(N)} \quad (1.97)$$

and the definition of the quantum volume can be further refined as follows

$$V_Q = \max_{n \leq N} \left(\min \left[n, \frac{1}{n\epsilon_{\text{eff}}(n)} \right]^2 \right) \quad (1.98)$$

Here the maximum is taken over all possible subsets of n qubits and keeps into account the fact that, according to the definitions above, the optimal circuit depth d formally decreases with N if more qubits with constant $\epsilon_{\text{eff}}(N)$ are added, and that the maximum volume could then be accessible to a given hardware by using only a portion of the available quantum register. As with qubit coherence times, an analog of Moore's law has also been proposed for the quantum volume: for example, IBM Research reported in the period 2017-2019 a yearly doubling of the maximal quantum volume achieved with their IBM Q prototypes, from $V_Q \simeq 4$ to $V_Q \simeq 16$ [177].

The challenges that must be faced to increase the available quantum volume are, by definition, closely related to the evolution of current NISQ devices into useful quantum simulators. On one hand, despite a few non-trivial issues related e.g. to individual qubit manipulation and frequency crowding (on superconducting platforms) or to finding new strategies for qubit confinement and spatial arrangement (on trapped ions processors), adding more qubits to the actual hardware mostly appears to be a purely technological challenge. On the other hand, implementing long sequences of high-fidelity operations on several qubits could be more demanding and may bring up some fundamental physical issues. Indeed, this requires considerable improvement of the average gate performances on multi-qubit devices, the suppression of both qubit decoherence and coherent errors due to imperfect qubit manipulations, and the reduction of unwanted qubit-qubit interactions (cross-talks) whose harmful effect increases with the system size. Along these lines, strategies to engineer better gate parallelization [178] and global entangling interactions [179] on trapped ions architectures have lately been put forward.

To conclude, and in addition to what was already mentioned in the previous sections, we also summarize here some of the theoretical contributions which in recent years aimed at expanding the range of applicability and the performances of digital quantum simulation protocols, independently of the specific hardware platform. Examples go from quantum chemistry and material sciences [57, 119, 180] to quantum field theories [181, 182] and high energy physics [183], from the dynamics of open [184–189] and many-body quantum systems [190] to linear response [63] and imaginary-time evolution [191, 192]. Refined product formulas for lattice simulation [193], genetic algorithms [194] and variational time evolution approaches directly inspired by the action minimization principle [195] have also been proposed in an attempt to increase the digital accuracy and resilience to hardware noise of quantum simulations algorithms.

Quantum dynamics on near-term quantum processors

In this chapter, we present original results for the digital quantum simulation of quantum dynamics on NISQ processors based on superconducting technology. First, some of the algorithms outlined in Chapter 1 will be put into practice with the aim of assessing the current capabilities of quantum processors. We will then move on beyond state-of-the-art results and present proof-of-principle demonstrations of new procedures and simulation techniques.

The quantum processors used in this work are provided by IBM Research within the IBM Quantum Experience (IBM Q) program. Results were obtained both on public free-access prototypes and on near-to-commercial chips. The latter were operated in close collaboration with the quantum computing research group at the IBM Zurich Laboratory.

The chapter is organized in three main sections: in the first part, we will provide some technical details about the quantum processors used in the experimental tests and the way they were programmed and operated on cloud. In the second part, we will present algorithms and results for the digital quantum simulation of some paradigmatic few-spin models, specifically adapted and designed to be run on IBM Q processors. In the third and most important part, we will then describe in detail the main theoretical and experimental achievements, namely the full digital quantum simulation of spin dynamical correlation functions and a thorough original discussion of device noise, error mitigation techniques and scalability perspectives. These findings will then be applied, with an original hybrid quantum-classical procedure, to extract the cross section of 4-dimensional inelastic neutron scattering processes on molecular magnetic systems.

2.1 The IBM Quantum Experience

IBM Q quantum processors were first made available for free cloud access in 2016. The physical chips are assembled and operated in the IBM Research Laboratory in Yorktown, while users can submit quantum circuits to be run by using graphical interfaces and the Qiskit python development kit [196]. At its heart, the latter manipulates QASM (Quantum Assembly Language) listings for quantum circuits [197]. Both the user interface and the control software are currently under active development and have undergone some deep restyling steps over the months in which the data presented here were taken, approximately January 2018-June 2019. However, the average properties and performances of the quantum processors, which we are going to outline below, were kept reasonably stable through constant maintenance and calibration. The IBM Quantum Experience can be accessed at <https://quantum-computing.ibm.com>, while Qiskit documentation and updates are hosted at <https://qiskit.org/>.

The processors

All IBM Q processors are based on superconducting technology and belong to the class of Noisy Intermediate Scale Quantum prototypes, having $N \leq 20$ physical qubits and featuring no quantum error correction. Individual qubits are realized as transmon circuits kept at a fixed frequency

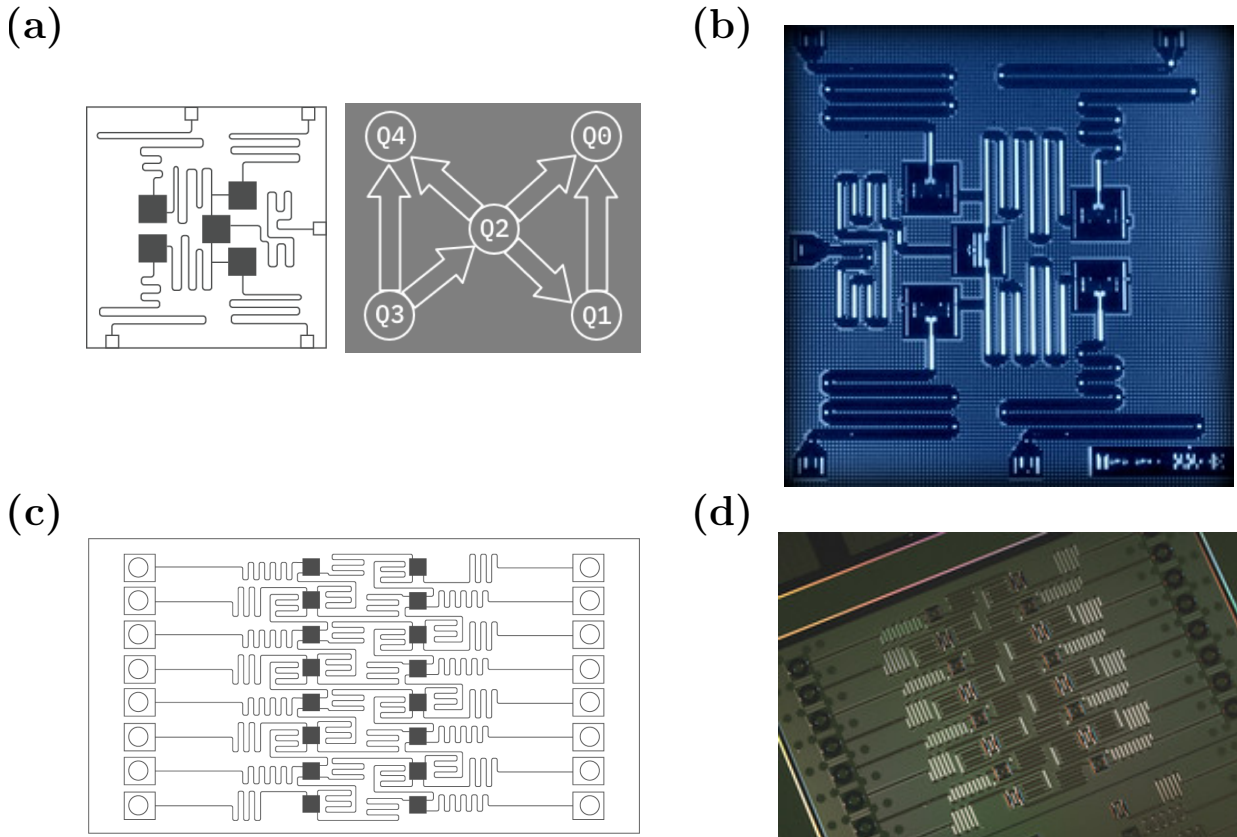


Figure 2.1: IBM Q public devices. (a) Layout of the 5-qubit “ibmqx4-Tenerife” processor, with a scheme of inter-qubit connectivity and CNOT gate directions. Black squares represent transmon qubits, coupled through coplanar waveguide resonators and controlled via transmission lines. (b) Optical image of a real IBM Q 5-qubit quantum processor. (c) Layout of the 16-qubit “ibmqx5-Rüschlikon” quantum processor. (d) Optical picture of a real IBM Q 16-qubit device. Images are reproduced under Creative Commons License, © IBM Research.

within the microwave range. Qubit-qubit couplings are obtained via coplanar waveguide (CPW) resonators, also named quantum bus cavities [158], and individual qubit control and readout is performed through a set of other CPW resonators put in contact with external transmission lines at the edge of the chip board [94, 166]. The processors are shielded from external electromagnetic influences and kept at cryogenic temperatures (typically 15-20 mK) inside dilution fridges. All data and results described in this chapter were obtained, via remote access, on 5-qubit (named “ibmqx2-Yorktown” and “ibmqx4-Tenerife”), 16-qubit (“ibmqx5-Rüschlikon” and “ibmq_16_melbourne”) and 20-qubit (“ibmq_20_tokyo”) processors. In Fig. 2.1 we show a schematic layout of typical quantum chip architecture and optical images of the real processors.

The fundamental set of quantum logic gates that is natively implemented on IBM Q processors [101] is essentially the \mathcal{S}_1 set described in Sec. 1.1.3, i.e. it contains single qubit rotations and CNOT. Microwave control pulses manipulating single qubits and activating two-qubit CNOT operations, shaped with gaussian and gaussian-derivative profiles, are optimized to minimize gate inaccuracies and population leakage outside the computational basis states of transmons [198, 199]. In terms of pulse optimization, it is worth mentioning that *virtual* z -rotations (i.e. phase gates) are used [100], see also Sec. 1.1.3: under this strategy, every single-qubit relative phase induced by a $R_z(\theta)$ operation is taken into account only as a change of reference frame in the control settings of all subsequent gates, without the need for physical manipulation. Virtual z -rotations are in principle error-free and instantaneous, such that better performances should be expected for circuits relying on the largest possible number of such quantum gates. A scheme of the four fundamental control sequences for IBM Q devices, to which all other user-requested quantum gates are reduced at the compilation stage, are reported in Fig. 2.2a.




Figure 2.2: Fundamental quantum operations on IBM Q devices. (a) Pulse scheme for the fundamental gate set. Every pulse is applied at the frequency reported at the beginning of each line: ω_q is the frequency of the qubit on which the U_i gates are applied, while ω_C and ω_T in the CNOT are the control and target qubit frequencies respectively. Image retrieved from <https://github.com/Qiskit/ibmq-device-information>. (b) Scheme of two transmon qubits coupled through an intermediate cavity resonator. Cross resonance interactions are activated by driving one qubit at the renormalized frequency of the other.

Single qubit operations. Quantum logic gates addressing individual qubits are obtained in three versions of different complexity and generality: at the lowest level, the quantum gate called $U_1(\lambda)$ is nothing but a virtual $\Phi(\lambda)$ phase gate, see Eq. (1.27)

$$U_1(\lambda) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\lambda} \end{pmatrix} = e^{i\frac{\lambda}{2}} R_z(\lambda) \quad (2.1)$$

and is obtained with just a Frame Change (FC) at the control software level, as commented above. A slightly more powerful gate is

$$U_2(\phi, \lambda) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -e^{i\lambda} \\ e^{i\phi} & e^{i(\lambda+\phi)} \end{pmatrix} \quad (2.2)$$

which requires one Gaussian Derivative (GD) and two FC pulses. Hadamard gates can for example be translated as $H = U_2(0, \pi)$. Finally, the most general $SU(2)$ gate, already introduced in Eq. (1.25), here is called U_3

$$U_3(\theta, \phi, \lambda) = \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i(\lambda+\phi)} \cos(\theta/2) \end{pmatrix} \quad (2.3)$$

and is implemented with two GD and three FC pulses at the frequency of the target qubit. Typical gating times for IBM Q devices are in the range of 50-150 ns for each GD pulse. It is also worth mentioning that, since the continuous phases are encoded through virtual frame changes, the total duration of single qubit rotations is fixed and does not depend on the choice of the θ , ϕ or λ parameters introduced above.

Two-qubit operations. The native two-qubit CNOT gate which completes the universal set of quantum operations is obtained with a cross-resonance (CR) scheme [104, 200] in which the control qubit is driven at the frequency of the target with a Gaussian Flattop (GF) pulse. The average GF duration is around a few hundreds of nanoseconds, typically $t_{GF} \simeq 100$ -400 ns. The physical basis of the cross-resonance gate can be understood starting from the circuit in Fig. 2.2b. Indeed, the typical Hamiltonian for a pair of qubits (Q1 and Q2, with frequencies ω_1 and ω_2 respectively) dispersively coupled through a common resonator reads [104]

$$H = \frac{\omega_1}{2} \sigma_z^{(1)} + \frac{\omega_2}{2} \sigma_z^{(2)} + J \sigma_x^{(1)} \sigma_x^{(2)} \quad (2.4)$$

Table 2.1: Typical average parameters for IBM Q devices

	ibmqx2	ibmqx4	ibmqx5	ibmq_16_melbourne	ibmq_20_tokyo
N of qubits	5	5	16	14	20
Frequency (GHz)	5.19	5.29	5.14	4.97	5.16
T_1 (μ s)	62.9	45.2	52.1	54.6	86.1
T_2 (μ s)	49.7	22.5	87.8	66.5	60.3
1-qubit gate error	$2 \cdot 10^{-3}$	$2 \cdot 10^{-3}$	$2 \cdot 10^{-3}$	$6 \cdot 10^{-3}$	$2 \cdot 10^{-3}$
2-qubit gate error	$5 \cdot 10^{-2}$	$4 \cdot 10^{-2}$	$5 \cdot 10^{-2}$	$5 \cdot 10^{-2}$	$5 \cdot 10^{-2}$
Readout error	$4 \cdot 10^{-2}$	$1 \cdot 10^{-1}$	$7 \cdot 10^{-2}$	$6 \cdot 10^{-2}$	$9 \cdot 10^{-2}$

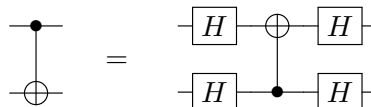
which upon diagonalization in the regime $J \ll \Delta_{12} = \omega_1 - \omega_2$ yields a set of renormalized qubit frequencies

$$\tilde{\omega}_1 = \omega_1 + \frac{J}{\Delta_{12}} \quad \tilde{\omega}_2 = \omega_2 - \frac{J}{\Delta_{12}} \quad (2.5)$$

In the new qubit basis, a drive Hamiltonian applied on Q1 at the frequency $\tilde{\omega}_2$ has the form [104]

$$H_{\text{drive}}(t) = A(t) \cos(\tilde{\omega}_2 t) \left(\sigma_x^{(1)} - \frac{J}{\Delta_{12}} \sigma_z^{(1)} \sigma_x^{(2)} + m_{12} \sigma_x^{(2)} \right) \quad (2.6)$$

where $A(t)$ is the amplitude shape of the pulse. Apart from single qubit effects $\sigma_x^{(1)}$ and spurious cross-talks $m_{12} \sigma_x^{(2)}$, the drive $H_{\text{drive}}(t)$ contains the generator $\sigma_z^{(1)} \sigma_x^{(2)}$ which can be tuned to obtain a CNOT gate. Without commenting further on these technical details, we mention that the whole treatment can be done symmetrically for a driving applied on Q2, and in general the whole 2-qubit SU(4) Lie algebra can be covered [102]. Quite relevant for the IBM Q realization is also the fact that, being Q1 and Q2 non-degenerate, when cross talks are taken into account [104] the effective strength of the cross resonance interaction is non-symmetric and is typically more favorable (requiring e.g. shorter gating times) for one particular choice of the control-target roles. As a consequence, and most often in public IBM Q devices, these relationships are fixed and are available to the user only in a specific configuration: for example, in Fig. 2.1a the directions of the arrows point from the designated control towards the target. Notice that these limitations can always be compensated through single qubit rotations, as in the following identity:



$$\text{CNOT}_{12} = H_1 H_2 \text{CNOT}_{21} H_1 H_2 \quad (2.7)$$

Qubit readout. Information about the state of individual qubits can be retrieved via measurements in the computational basis. Each quantum circuit run on IBM Q devices is repeated a number $n_{\text{shots}} \leq 8192$ of times, from which the statistics of the outcomes can be reconstructed, leading to the probabilities p_0 and p_1 of measuring a qubit in states $|0\rangle$ and $|1\rangle$ respectively. At the hardware level, the state of each qubit is measured by sending a microwave pulse into readout resonators: the reflected signals are then amplified via a Josephson parametric converter followed by high electron mobility transistor amplifiers operating at 4 K.

We conclude this section by providing in Tab. 2.1 a set of typical physical characterization parameters for the IBM Q quantum processors. These are averaged over all qubits and a significant variability of qubit quality must usually be taken into account. Most of these properties are updated on a daily basis through hardware recalibration. Additional information and technical details are maintained at <https://github.com/Qiskit/ibmq-device-information>.

Writing quantum software

Remote user interaction with IBM Q devices and quantum programming happen via the open source python development kit called Qiskit [196]. During the time frame in which the results presented here were obtained, the Qiskit distribution went through several releases, moving from

version 0.4 to 0.10. Although the overall software interface gradually evolved and has seen, among other improvements, the addition of field-specific libraries and a better handling of backends and circuit compilation routines, the underlying structure for writing a sequence of quantum gates and running it on a real quantum processor has remained essentially the same. Without going into details, we can give the flavor of a Qiskit program with the following lines of code, where we create an instance of a 2-qubit quantum register, we perform a simple sequence of gates which brings a factorized $|00\rangle$ default initial state into the Bell state $\sqrt{2}|\Phi^+\rangle = |00\rangle + |11\rangle$ and then measures both qubits in the computational basis, writing the outcomes on a 2-bit classical register:


```
import qiskit as qk

qr = qk.QuantumRegister(2, name='qr')
cr = qk.ClassicalRegister(2, name='cr')
qc = qk.QuantumCircuit(qr, cr)

qc.h(qr[0])
qc.cx(qr[0], qr[1])

qc.measure(qr[0], cr[0])
qc.measure(qr[1], cr[1])
```

The corresponding circuit representation is



When parsed and compiled to be run on a real IBM Q processor, the above program is converted into an OpenQASM listing, featuring the native IBM Q gates. Remembering that for the Hadamard gate we can write $H = U_2(0, \pi)$, it is straightforward to verify that the corresponding quantum assembly script is

```
OPENQASM 2.0;
include "qelib1.inc";
qreg q[5];
creg cr[2];
u2(0.0,3.14159265358979) q[1];
cx q[1],q[0];
u2(0.0,3.14159265358979) q[0];
u2(0.0,3.14159265358979) q[1];
barrier q[1],q[0];
measure q[0] -> cr[0];
measure q[1] -> cr[1];
```

In this case, the target backend was the “ibmqx4-Tenerife” processor, and the compiler used the identity of Eq. (2.7) to account for the fixed control-target CNOT relationship between physical qubits Q1 and Q0 imposed by the chip, see Fig. 2.1a. This ability of the Qiskit compiler to adapt a given circuit to a certain hardware, for example by taking into account the actual qubit-qubit connectivity and CNOT directionality, has greatly improved over time. However, in order to maintain full control over this mapping and to avoid unexpected modifications of the quantum circuit, for all results presented here the adaptation to the real backend topology was always hard-

coded, including e.g. SWAP operations

$$\text{SWAP} = \begin{array}{c} \bullet \quad \oplus \quad \bullet \\ | \quad | \quad | \\ \oplus \quad \bullet \quad \oplus \end{array} \quad (2.8)$$

explicitly when needed. This was of course possible due to the small size of the quantum registers involved. Moreover, the choice of the subset of qubits to be used to run a certain algorithm on a given processor was also requested explicitly, in such a way that different realizations could be compared and the best qubits in terms of coherence and gate fidelity could be targeted. It is worth noting that, along with Qiskit, the ecosystem of open source quantum software libraries is already quite well developed [201], with both full-stack approaches (e.g. IBM Q with Qiskit, Rigetti Computing with PyQuil and Forest¹) and more application-oriented efforts.

2.2 Digital quantum simulations of spin Hamiltonians

Initial experiments on IBM Q processors were performed primarily as a way of assessing the bare performances of the hardware and to establish a working library of gate sequences. In this section, we present early results on the digital quantum simulation of some paradigmatic spin models. Despite an overall quantitative accuracy which often turns out to be not fully satisfying, these preliminary results will help to identify a set of features and physical quantities of interest that the prototype quantum processors were already able to reproduce. Leveraging this knowledge of the capabilities and intrinsic limitations of IBM Q processors will turn out to be crucial in the design and successful completion of more complex studies, which will be presented afterwards in Sec. 2.3. All simulations reported in this section were performed in March 2018, and some of the results were later published in Ref. [202].

2.2.1 Tunneling of $S = 1$ total magnetization

The magnetic moment of a spin $S = 1$ particle under the action of single-spin crystal-field anisotropy obeys a Hamiltonian of the form [47, 203]

$$\mathcal{H} = DS_z^2 + E(S_x^2 - S_y^2) \quad (2.9)$$

The mapping of this model to a qubit register can be made straightforwardly by considering, on one hand, the close relationship between qubits and spin $s = 1/2$ particles, and on the other hand the fact that the $S = 1$ components can be seen as the sum of two spin-1/2 operators:

$$S_\alpha = s_\alpha^{(1)} + s_\alpha^{(2)} \quad \text{for } \alpha = x, y, z \quad (2.10)$$

The corresponding mapped Hamiltonian then becomes [47, 54]

$$H = 2Ds_z^{(1)}s_z^{(2)} + 2E \left(s_x^{(1)}s_x^{(2)} - s_y^{(1)}s_y^{(2)} \right) \quad (2.11)$$

Here, in view of studying the time evolution of the total magnetization $S_z = s_z^{(1)} + s_z^{(2)}$, we have neglected operators corresponding to constants of the motion commuting with $s_z^{(i)}$. By using the relation $s_\alpha^{(i)} = (1/2)\sigma_\alpha^{(i)}$, the Hamiltonian H can then be written in Pauli form

$$H = \frac{D}{2}\sigma_z^{(1)}\sigma_z^{(2)} + \frac{E}{2} \left(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)} \right) \quad (2.12)$$

The time evolution of the expectation value of S_z then reads

$$\begin{aligned} \langle S_z(t) \rangle &= \langle \psi_0 | e^{iHt} \left(s_z^{(1)} + s_z^{(2)} \right) e^{-iHt} | \psi_0 \rangle \\ &= \langle \psi_0 | e^{i\frac{Et}{2}(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)})} \left(s_z^{(1)} + s_z^{(2)} \right) e^{-i\frac{Et}{2}(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)})} | \psi_0 \rangle \end{aligned} \quad (2.13)$$

¹See for example <https://github.com/rigetti/pyquil> and <https://www.rigetti.com/forest>.




Figure 2.3: Digital quantum simulation of $S = 1$ magnetization on an IBM Q processor. (a) Scheme and connectivity of the IBM Q quantum processor ibmqx2-Yorktown. Retrieved from <https://github.com/Qiskit/ibmq-device-information>. (b) Results of the 2-qubit digital quantum simulation using the quantum circuit in Eq. (2.17). Data are shown for two different choices of qubits on the quantum processor. In both cases, the CNOT direction was chosen according to the device characteristics and the statistics of the occupation probabilities for $s_z^{(i)}$ eigenstates was reconstructed by running the same sequence of quantum gates $n_{\text{shots}} = 8192$ times, i.e. the maximum allowed on IBM Q public devices.

where $|\psi_0\rangle$ is the initial state of the system and we have used the relation

$$[\sigma_\alpha^{(1)}\sigma_\alpha^{(2)}, \sigma_\beta^{(1)}\sigma_\beta^{(2)}] = 0 \quad \text{for } \alpha, \beta = x, y, z \quad (2.14)$$

As a result, it is clear that non-trivial evolution is only generated by the term

$$H_E = \frac{E}{2} \left(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)} \right) \quad (2.15)$$

which we can readily compute on a quantum register. In fact, this is probably one of the simplest possible examples of digital quantum simulation, as the time evolution operator

$$U(t) = e^{-i\frac{Et}{2}(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)})} = e^{-i\frac{Et}{2}\sigma_x^{(1)}\sigma_x^{(2)}} e^{i\frac{Et}{2}\sigma_y^{(1)}\sigma_y^{(2)}} \quad (2.16)$$

does not require digital approximations through the ST formula of Eq. (1.7). Within the IBM Q native set of operations, a quantum circuit achieving the desired result for a target phase evolution $\lambda = Et/2$ is for example

$$\begin{array}{c} \text{---} X \text{---} U_3^x(\pi/2) \text{---} \bullet \text{---} U_3^x(-\pi/2) \text{---} U_3^y(-\pi/2) \text{---} \bullet \text{---} U_3^y(\pi/2) \text{---} \\ \text{---} X \text{---} U_3^x(\pi/2) \text{---} \oplus \text{---} U_1(-2\lambda) \text{---} \oplus \text{---} U_3^x(-\pi/2) \text{---} U_3^y(-\pi/2) \text{---} \oplus \text{---} U_1(2\lambda) \text{---} \oplus \text{---} U_3^y(\pi/2) \text{---} \end{array} \quad (2.17)$$

where $U_3^x(\theta) = U_3(\theta, -\pi/2, \pi/2) = R_x(\theta)$, $U_3^y(\theta) = U_3(\theta, 0, 0) = R_y(\theta)$ and we have used quantum gate decompositions similar to the ones in Eqs. (1.48)-(1.52). The X gates at the beginning prepare the initial state $|11\rangle$ from the default $|00\rangle$ assumed in Qiskit. The elementary digital quantum simulation algorithm in Eq. (2.17) was tested on different pairs of qubits on the ibmqx2 quantum processor, as reported in Fig. 2.3. Data are extracted by measuring the quantum register in the computational basis, i.e. by reconstructing occupation probabilities for the eigenstates of $s_z^{(i)} = (1/2)\sigma_z^{(i)}$. We notice a very good agreement with theoretical predictions featuring the tunneling oscillations of the total magnetization, with only minor discrepancies with respect to ideal results. Such differences come essentially in the form of a slight damping of the oscillations, whereas the oscillation frequency is correctly reproduced. In fact, we can already mention that this is a typical effect of the combined action of noise channels and systematic errors on the quantum chip, as will become more clear in the following.




Figure 2.4: Digital quantum simulation of spin-1/2 chains on IBM Q processors. (a) Spin magnetization at one end and at the center of a 3-qubit Heisenberg chain, simulated on ibmqx2 for one digital step and with initial state $|\downarrow\uparrow\downarrow\rangle$. (b) Spin magnetization at one end and at the center of a 3-qubit chain with anti-symmetric exchange interactions, simulated on ibmqx5 for one digital step and with initial state $|\downarrow\uparrow\downarrow\rangle$.

2.2.2 Spin-1/2 chains

A slight increase in the complexity of the target digital simulation can be obtained by combining SU(4) Hamiltonians for different spin-1/2 pairs, thus building small-scale examples of spin chains. In this case, bonds with common ends typically generate non-commuting terms and thus call for a non-trivial use of the Suzuki-Trotter digitalization procedure. One of the simplest prototypes is a system of three spin-1/2 particles arranged in a linear chain topology with open ends and Heisenberg-like interactions. The Hamiltonian reads

$$H_{\text{Heis},3} = H_{\text{Heis},2}^{12} + H_{\text{Heis},2}^{23} \quad (2.18)$$

where

$$H_{\text{Heis},2}^{ij} = J_{ij} \left(\sigma_x^{(i)} \sigma_x^{(j)} + \sigma_y^{(i)} \sigma_y^{(j)} + \sigma_z^{(i)} \sigma_z^{(j)} \right) \quad (2.19)$$

It is straightforward to verify by direct inspection that $[H_{\text{Heis},2}^{12}, H_{\text{Heis},2}^{23}] \neq 0$. The digital quantum simulation is therefore built as

$$e^{-iH_{\text{Heis},3}t} \simeq \left(e^{-iH_{\text{Heis},2}^{23}t/n} e^{-iH_{\text{Heis},2}^{12}t/n} \right)^n \quad (2.20)$$

and a single digital step corresponds to a quantum circuit of the form

$$\begin{array}{c} \boxed{e^{-iH_{\text{Heis},2}^{12}t/n}} \\ \hline \boxed{e^{-iH_{\text{Heis},2}^{23}t/n}} \end{array} \quad (2.21)$$

We take for simplicity $J_{12} = J_{23} = J$ and we define $\lambda = Jt$ and $\lambda_n = \lambda/n$. We can therefore decompose each 2-qubit block in Eq. (2.21) with the usual techniques in CNOT gates and single qubit rotations (see for example Fig. 1.4a). Using IBM Q native gates a typical sequence for $e^{-iH_{\text{Heis},2}^{ij}t/n}$ is then

$$\begin{array}{c} \boxed{\text{ZZ}(\lambda_n)} \quad \boxed{U_3^x(\pi/2)} \quad \boxed{\text{ZZ}(\lambda_n)} \quad \boxed{U_3^x(-\pi/2)} \quad \boxed{U_3^y(-\pi/2)} \quad \boxed{\text{ZZ}(\lambda_n)} \quad \boxed{U_3^y(\pi/2)} \\ \boxed{U_3^x(\pi/2)} \quad \boxed{U_3^x(-\pi/2)} \quad \boxed{U_3^y(-\pi/2)} \quad \boxed{U_3^y(\pi/2)} \end{array} \quad (2.22)$$

where

$$\boxed{\text{ZZ}(\lambda_n)} = \begin{array}{c} \bullet \\ \oplus \\ \boxed{U_1(2\lambda_n)} \\ \oplus \\ \bullet \end{array} \quad (2.23)$$

Here, in the spirit of testing the hardware capabilities and resilience, we do not use any CNOT reduction scheme (such as e.g. the one shown in Fig. 1.4b), in such a way that each 2-qubit Heisenberg block consists of 6 CNOT gates and single-qubit rotations. A single digital step on three qubits has therefore a circuit depth corresponding to 12 CNOTs. Results obtained for $n = 1$ on the ibmqx2 quantum processor are shown in Fig. 2.4a. While the overall oscillation of the spin magnetization is visible, it is clear that the quantum hardware is pushed here close to its present limitations in terms of quantitative accuracy. In particular, $n > 1$ digital steps are still beyond reach on this particular implementation of the Heisenberg model: in terms of quantum gates, we empirically find approximately 10-15 CNOTs in the circuit to be an upper bound for reliable results. Consistent performances on similar SU(4) Hamiltonians are also reported in the recent literature [62].

Another series of experiments, whose results are shown in Fig. 2.4b, targeted a 3-qubit chain with anti-symmetric exchange interactions (ASE)

$$\mathbf{H}_{\text{ASE},3} = \mathbf{H}_{\text{ASE},2}^{12} + \mathbf{H}_{\text{ASE},2}^{23} \quad (2.24)$$

where

$$\mathbf{H}_{\text{ASE},2}^{ij} = J \left(\sigma_x^{(i)} \sigma_y^{(j)} - \sigma_y^{(i)} \sigma_x^{(j)} \right) \quad (2.25)$$

Here, a single 2-qubit block $e^{-i\mathbf{H}_{\text{ASE},2}^{ij}t/n}$ within a digital step takes the form

$$\begin{array}{c} \boxed{U_3^y(\pi/2)} \quad \boxed{\text{ZZ}(\lambda_n)} \quad \boxed{U_3^y(-\pi/2)} \quad \boxed{U_3^x(-\pi/2)} \quad \boxed{\text{ZZ}(-\lambda_n)} \quad \boxed{U_3^x(\pi/2)} \\ \boxed{U_3^x(\pi/2)} \quad \boxed{U_3^x(-\pi/2)} \quad \boxed{U_3^y(-\pi/2)} \quad \boxed{U_3^y(\pi/2)} \end{array} \quad (2.26)$$

and the total 3-qubit quantum algorithm

$$\begin{array}{c} \boxed{e^{-iH_{\text{ASE},2}^{12}t/n}} \\ \hline \boxed{e^{-iH_{\text{ASE},2}^{23}t/n}} \end{array} \quad (2.27)$$

contains 8 CNOTs per digital step. As it can be seen in the experimental results obtained on the ibmqx5 quantum processor, oscillations are once again accurately reproduced, while the total amplitude is slightly damped. In both ASE and Heisenberg chains, it is worth noticing that a source of quantitative inaccuracy is also an apparent shift of the oscillation maxima with respect




Figure 2.5: Digital quantum simulation of the Transverse Field Ising model. (a) Quantum circuit for a single Trotter digital step, using IBM Q native quantum logic gates. (b)-(d) Time evolution of the total spin magnetization along the z direction for increasing number n of digital steps. Exact Hamiltonian diagonalization (dashed line) and numerical matrix exponentiation using the Trotter formula (solid blue line) are compared with experimental results obtained on the ibmqx5 quantum processor.

to the exact results. While in a fair scenario in which classical simulations are not available this discrepancy cannot in general be detected, it will be nevertheless interesting to investigate it further in the following. Indeed, we will see in later sections that this issue is rather ubiquitous in many of the experimental results obtained on real hardware and, at least to some extent, it reflects systematic inaccuracies in the implementation of quantum gates. Under this light, we will be able in some special cases to exploit the underlying physical properties and symmetries of the target models to mitigate their impact *a posteriori* without resorting to explicit exact solutions.

2.2.3 Transverse Field Ising model

We conclude this exploration of the basic performances of IBM Q processors as digital quantum simulators by showing on an explicit example the effect of different choices for the number of Suzuki-Trotter approximation steps. As we extensively discussed in Chapter 1, a dense discretization is required to accurately reproduce the physical reality of a target model. However, this in turn means that long enough quantum computations must be implemented and carried out: on NISQ processors, there is inevitably a trade-off between digital error and experimental limitations. Here we use as a test case the 2-qubit Transverse Field Ising model, introduced in Eq. (1.91)

$$H_{\text{TIM},2} = \frac{Bg}{2} (\sigma_x^{(1)} + \sigma_x^{(2)}) + J_{zz} \sigma_z^{(1)} \sigma_z^{(2)} \quad (2.28)$$

The corresponding quantum circuit for a digital step, reported in Fig. 2.5a (see also Eq. (1.93) for the corresponding mathematical formulation) is sufficiently simple to make a few Trotter steps well within the current boundaries of experimental feasibility, even if with lower and lower accuracy as n increases. The cases with $n = 1, 5, 10$ simulated on ibmqx5 are reported in Fig. 2.5b-d, where experimental results for the total spin magnetization in the case $Bg = 2J_{zz}$ are compared to the numerical solution of the TIM both in the exact and digital version. The number of digital steps is kept constant over the whole range of the simulation. The quantitative and qualitative agreement is good up to approximately $n \simeq 5$ (i.e. 10 CNOTs) and $J_{zz}t \simeq 18$, while it is also clear the at least $n \simeq 10$ is needed to start getting physically meaningful results at long times.

The lesson that must be learned from all the results presented in this section is that there is still much room for improvement in all those applications in which the time evolution of physical quantities is directly the object of the simulation. However, as we will more extensively discuss in the second part of this chapter, it is also true that there exist a set of questions that can be asked about a physical system, starting from e.g. normal oscillation frequencies, for which a useful answer can already come from digital quantum simulation algorithms run on NISQ processors.

2.3 Dynamical quantum correlations

It is well known that a good and reliable estimation of quantum mechanical normal frequencies is critical for exploring the energetic structure of a target system. The most general physical quantities containing such information are dynamical correlation functions [122, 123, 125], which were already briefly introduced in Sec. 1.1.5. Indeed, dynamical correlation functions describe how excitations, possibly induced by external probes, propagate inside a system and are thus common to all kinds of investigations involving e.g. linear response, energy transport and scattering. The results presented in the previous section suggest that an analysis of dynamical correlations and of the energetic and structural information they carry, if faced with the necessary contributions from data and error analysis, might be within reach of a digital quantum simulation scheme performed on NISQ processors. As dynamical correlations for quantum systems are typically hard to obtain numerically on classical devices, the successful use of a quantum hardware might open new and interesting perspectives in the field.

Here we will present a systematic procedure to compute and analyze dynamical correlation functions on spin systems, and we will provide proof-of-principle realizations obtained on IBM Q processors. Besides the intrinsic theoretical and experimental value of such results, the final goal will be the study of real physical systems, such as magnetic clusters made of atomic or molecular spins: indeed, we will demonstrate how quantum dynamical correlations can be used to efficiently extract the cross section for an investigation technique known as 4-dimensional inelastic neutron scattering, which nicely relates the general knowledge of paradigmatic spin models to a specific real world application. During the discussion, we will have the opportunity to characterize in greater detail the major sources of experimental error on the real processors, starting from an almost completely hardware-agnostic end user perspective. The most relevant results in this section appeared in Ref. [124].

2.3.1 Quantum algorithms for dynamical correlation functions

The general procedure for computing and extracting dynamical correlation functions by making use of digital quantum simulation techniques on a register of qubits was introduced and described in Sec. 1.1.5. Here we apply that scheme to some 2- and 3-spin systems, which we will also call molecules from now on, adapting the required gate sequences to IBM Q devices. All the benchmark spin models considered here share the general structure

$$H = \sum_{\substack{i=1 \\ \alpha=x,y,z}}^{N-1} J_i^\alpha s_\alpha^{(i)} s_\alpha^{(i+1)} + B \sum_{i=1}^N g_i s_z^{(i)} \quad (2.29)$$

Table 2.2: Parameters for simulated spin molecules. For each molecule, the parameter J sets the overall energy scale.

	molecule 1	molecule 2	molecule 3	molecule 4	molecule 5
N of spins	2	2	2	3	3
Bg_1	$3J$	$10J$	$10J$	$10J$	$10J$
Bg_2	$3J$	$12.5J$	$12.5J$	$10J$	$10J$
Bg_3	-	-	-	$10J$	$10J$
J_1^α	J	J	$J_1^z = J, J_1^{x,y} = 0$	J	$J_1^x = J_1^y = -0.5J_1^z = J$
J_2^α	-	-	-	J	$J_1^x = J_1^y = -0.5J_1^z = 0.7J$

where $s_\alpha^{(i)} = (1/2)\sigma_\alpha^{(i)}$ are spin-1/2 operators. The target spin-spin dynamical correlation functions are defined as

$$\mathcal{C}_{ij}^{\alpha\beta}(t) = \langle 0 | s_\alpha^{(i)}(t) s_\beta^{(j)} | 0 \rangle = \sum_p \langle 0 | s_\alpha^{(i)} | p \rangle \langle p | s_\beta^{(j)} | 0 \rangle e^{-iE_p t} \quad (2.30)$$

where $|0\rangle$ is the energetic ground state, $|p\rangle$ are molecular eigenstates and E_p are the corresponding eigenenergies. We notice that, already from the expansion above, it is immediately clear that a full knowledge of dynamical correlations is essentially equivalent to a complete characterization of the system under study, i.e. a diagonalization of the relevant Hamiltonian.

For the $N = 2$ case, we consider three different choices of the structural parameters, corresponding to different degrees of complexity both in the required simulation algorithm and in the resulting dynamical behavior. The parameters are collected in Tab. 2.2: molecule 1 is a Heisenberg dimer of equivalent spins ($g_1 = g_2 = g$) put in an external field, molecule 2 has Heisenberg interactions and inequivalent spins ($g_1 \neq g_2$), and molecule 3 features inequivalent spins and a simple Ising term aligned to the external field. Only molecule 2 requires a non-trivial application of the Trotter approximation formula. On the other hand, the $N = 3$ case always requires a digital Trotter discretization due to the non-commutativity of the Hamiltonian interaction terms corresponding to edges with common ends, just like it was the case in Sec. 2.2.2. Two different spin trimers were considered, as reported again in Tab. 2.2: molecule 4 is a fully isotropic Heisenberg system in an external field, while molecule 5 features a anisotropic XYZ interaction with inequivalent bonds. In all cases, the external field was chosen to be sufficiently strong so that the ground state could *a priori* be assumed to be approximately $|0\rangle = |\downarrow\downarrow\rangle$ ($|\downarrow\downarrow\downarrow\rangle$ for trimers). Following the usual conventions adopted in quantum computation [5], we make the identification $|\downarrow\rangle = |1\rangle$ between spin-1/2 and qubit states. More in general, a separate analysis is required, possibly via variational approaches such as the VQE algorithm, to obtain the ground state and to prepare it on the quantum register.

The general structure of the quantum algorithm to extract dynamical correlation functions is reported in Fig. 2.6a for dimers and Fig. 2.6b for trimers. The final measure on the ancilla provides, as described in Sec. 1.1.5, real and imaginary parts of the correlations

$$\tilde{\mathcal{C}}_{ij}^{\alpha\beta}(t) = \langle 0 | \sigma_\alpha^{(i)}(t) \sigma_\beta^{(j)} | 0 \rangle \quad (2.31)$$

which correspond to $\mathcal{C}_{ij}^{\alpha\beta}(t)$ up to a constant 1/4 factor. In both dimer and trimer cases, the choice of the target qubits in the initial and final controlled operations determines the pair of spins on which the dynamical correlation is computed, possibly including autocorrelations. These controlled unitaries involving Pauli operations can be obtained within the IBM Q native set as follows

$$\begin{array}{c} \bullet \\ | \\ \square \alpha \end{array} \equiv \begin{array}{c} \bullet \\ | \\ \square \sigma_\alpha \end{array} = \begin{array}{c} \bullet \\ | \\ \square R \oplus \square R^\dagger \end{array} \quad (2.32)$$

where $R = H = U_2(0, \pi)$ for $\alpha = z$, $R = U_1(\pi/2)$ for $\alpha = y$ and no R is needed for $\alpha = x$. The time evolution terms $U_{ij}(t)$ are designed according to the usual digital quantum simulation procedures and applying Suzuki-Trotter digital approximation when required. In particular, for general 2-qubit




Figure 2.6: Digital quantum simulation of spin-spin dynamical correlations. (a) Quantum circuit for cross-correlations on a spin dimer (molecules 1, 2 and 3). Interaction with the external magnetic field is implicitly included in the $U_{12}(t)$ unitary operation. (b) Quantum circuit for autocorrelations on a spin trimer (molecules 4 and 5). Here the rotations induced by an external field are explicitly factored out as they commute ($g_i = g$ for $i = 1, 2, 3$) with the two-body interactions. The latter must be treated with Trotter digital steps, and $\tau = t/n$. In both panels, α, β are shorthand notations for $\sigma_{\alpha, \beta}$ and U_{meas} allows to retrieve from a final measurement in the computational basis the expectation values of σ_x (for $U_{\text{meas}} = H$) and σ_y (for $U_{\text{meas}} = R_x(\pi/2)$) on the ancilla a . In turn, these encode respectively the real and imaginary parts of the target correlation function.

SU(4) interactions leading to time evolution terms of the form

$$U_{ij}(t) = e^{-it \sum_{\alpha} J^{\alpha} s_{\alpha}^{(i)} s_{\alpha}^{(j)}} \quad (2.33)$$

we make use of the following decomposition which optimizes the circuit depth per digital step [127], see also Fig. 1.4b

$$e^{-it \sum_{\alpha} J^{\alpha} s_{\alpha}^{(i)} s_{\alpha}^{(j)}} = \begin{array}{c} \bullet \\ \oplus \\ \oplus \\ \oplus \end{array} \begin{array}{c} \boxed{R_x(\frac{J^x t}{2} - \frac{\pi}{2})} \\ \boxed{R_z(\frac{J^z t}{2})} \\ \boxed{R_z(\frac{-J^y t}{2})} \\ \boxed{R_x(\frac{\pi}{2})} \end{array} \begin{array}{c} \bullet \\ \oplus \\ \oplus \\ \oplus \end{array} \begin{array}{c} \boxed{H} \\ \boxed{R_x(\frac{\pi}{2})} \\ \boxed{H} \\ \boxed{R_x(\frac{\pi}{2})} \end{array} \quad (2.34)$$

For the case of molecule 4 and 5 the trotterization is explicitly shown in Fig. 2.6b. For the spin dimers, $n > 1$ digital steps are required only in the simulation of molecule 2, and here we make use of the second order Suzuki-Trotter formula introduced in Eq. (1.68). The time evolution part of the

circuit then becomes

$$\left(\begin{array}{c} \boxed{R_z(\frac{Bg_1\tau}{2})} \text{---} \bullet \text{---} \boxed{R_x(\frac{J^x\tau}{2} - \frac{\pi}{2})} \text{---} \boxed{H} \text{---} \bullet \text{---} \boxed{H} \text{---} \bullet \text{---} \boxed{R_x(\frac{\pi}{2})} \text{---} \boxed{R_z(\frac{Bg_1\tau}{2})} \\ \oplus \text{---} \boxed{R_z(\frac{Bg_2\tau}{2})} \text{---} \oplus \text{---} \boxed{R_z(\frac{J^z\tau}{2})} \text{---} \oplus \text{---} \boxed{R_z(-\frac{J^y\tau}{2})} \text{---} \oplus \text{---} \boxed{R_x^\dagger(\frac{\pi}{2})} \text{---} \boxed{R_z(\frac{Bg_2\tau}{2})} \end{array} \right)^n \quad (2.35)$$

where $\tau = t/n$. Being intrinsically limited by the hardware capabilities, for molecule 2 we use $n = 2$ up to $Jt = 2.0$ and $n = 4$ for $2 < Jt \leq 6.0$, and we set $n = 2$ for all evolution phases for molecules 4 and 5². While not optimal in terms of the computation of exact physical values for the dynamical correlations at long times, these choices still allow for a reliable extraction of the correct normal frequencies, as it will be demonstrated shortly by making use of suitable classical data analysis techniques. In this sense, the hybrid procedure we are discussing, while remaining relatively far from the full advantages promised by a post-NISQ digital quantum simulator, is nevertheless able to tackle real problems and provide physically relevant answers, holding promise for speed-ups over classical counterparts.

2.3.2 Results and error mitigation

From the structure of the algorithm presented in Fig. 2.6 it is straightforward to see that, in order to avoid the necessity for SWAP operations and thus some CNOT overhead, the role of the ancilla must be assigned on the quantum hardware in such a way that the chosen qubit can directly be involved in a controlled operation with the qubits appearing in the target dynamical correlation function to be computed. This means, for example, that the quantum circuit for dimers, Fig. 2.6a, requires at least a set of three all connected qubits. For this purpose, we used the bow-tie shaped ibmqx4 chip, as shown in Fig. 2.7a. On the other hand, the case of three-spin molecules, Fig. 2.6b, requires four qubits with, ideally, the ancilla put in direct contact with all other three qubits individually. On IBM Q public devices, this configuration is not available, and a hybrid solution must be adopted: as reported in Fig. 2.7a, the 5-qubit ibmqx4 can be used to compute autocorrelations (i.e. functions of the form $C_{ii}^{\alpha\beta}(t)$) and nearest-neighbors cross-correlations (e.g. $C_{21}^{\alpha\beta}(t)$ and $C_{32}^{\alpha\beta}(t)$), while the 16-qubit ibmqx5 can easily accommodate the quantum circuit for spin 1 and spin 3 autocorrelations and next-to-nearest neighbors (i.e. of the form $C_{31}^{\alpha\beta}(t)$) cross-correlations. The 20-qubit chip ibmq_20_tokyo, which was used for molecule 5, can simultaneously offer access to all dynamical correlations without any remapping.

We start presenting results on real IBM Q processors with the series of raw experimental data obtained through the IBM Quantum Experience on molecule 1 for $C_{ij}^{xx}(t)$ correlations, reported in Fig. 2.7b. As it can immediately be seen, the qualitative similarities with the ideal numerically computed solution are nevertheless strongly affected by a combination of damping and horizontal shift (i.e. phase error) effects. We recall that similar problems were encountered in Sec. 2.2. Leaving for a successive paragraph a more detailed analysis and justification of how these errors arise and are related to specific hardware limitations, we will now present a method for reducing the impact of the quantitative inaccuracy. This procedure, which we call Phase-and-Scale (PaS) error mitigation, exploits some mathematical properties of the dynamical correlation functions which can be inferred *a priori* without any knowledge of the exact solution of the simulation problem. Error mitigation techniques have recently emerged as useful tools to extend and improve the computational capabilities of noisy quantum hardware [62, 170–172]: it is however important to point out explicitly that error mitigation is intimately different from quantum error correction. On one hand, the latter generally involves the use of several physical qubits to encode and protect from external noise a much smaller set of logical qubits, in a configuration which is still inaccessible to the present few-qubit NISQ processors. On the other hand, error mitigation is a form of data post-processing which can polish and extrapolate experimental results obtained on quantum processors. Indeed, all error mitigation techniques can only be effective when the raw data retain some form of structure and are well distinct from random noise: as a result, this approach remains confined within the typical

²Further comments and plots comparing Trotter and exact evolutions are provided in Appendix A.




Figure 2.7: Digital quantum simulation of dynamical correlations on IBM Q devices. (a) Mapping of the spin dimers and trimers, along with the ancilla for readout of dynamical correlations, on *ibmqx4*, *ibmqx5* and *ibmq_20_tokyo* IBM Q devices. Qubit labels are given according to Fig. 2.6. Notice that the direction of some of the CNOTs on *ibmqx4* changed over successive device calibrations and that a straightforward remapping $1 \leftrightarrow 3$ of the trimer on *ibmqx4* can be used to put the ancilla in contact with qubits 2 and 3 instead of 1 and 2. (b) Raw data for molecule 1, computed on *ibmqx4* (circles) and compared to exact results (solid lines). (c) Phase-and-Scale error mitigated data for molecule 1 and corresponding fit. The extracted Fourier components are sufficiently good that there is no appreciable difference between the fit and the exact results. (d) PaS-mitigated results for molecule 2, computed on *ibmqx4*, and corresponding fit. The same color code as in panels (b)-(c) is used for real and imaginary parts. (e) Real part of dynamical correlations for molecule 4 and corresponding fit. Auto- and nearest neighbors cross-correlations are computed on *ibmqx4*, while next-to-nearest neighbors cross-correlations are computed on *ibmqx5*. Fitted Fourier coefficients and additional experimental data are reported in Appendix A. All results here were obtained by averaging the quantum circuit outcomes over $n_{\text{shots}} = 8192$.

coherence time of the qubits at use. In our case, we first observe that, for $t = 0$, autocorrelation functions of the form $\mathcal{C}_{ii}^{\alpha\alpha}(0)$ are a sum of positive squares of absolute values:

$$\mathcal{C}_{ii}^{\alpha\alpha}(0) = \sum_p |\langle 0 | s_\alpha^{(i)} | p \rangle|^2 \quad (2.36)$$

We can impose $\mathcal{C}_{ii}^{\alpha\alpha}(0)$ to be real and positive by taking the experimental points for the real and imaginary parts, which are obtained with different final measurements (see Fig. 2.6) and are typically both different from zero, computing a phase correction

$$\phi = \arg(\operatorname{Re}[\mathcal{C}_{ii}^{\alpha\alpha}(0)] + i \operatorname{Im}[\mathcal{C}_{ii}^{\alpha\alpha}(0)]) \quad (2.37)$$

and applying it to rotate the complex number back on the real axis

$$\mathcal{C}_{ii,\text{mitigated}}^{\alpha\alpha}(0) \equiv \hat{\mathcal{C}}_{ii}^{\alpha\alpha}(0) = e^{-i\phi} \mathcal{C}_{ii}^{\alpha\alpha}(0) \quad (2.38)$$

We heuristically find, by comparing to ideal numerical solutions, that the same phase correction ϕ computed at $t = 0$ by imposing that $\mathcal{C}_{ii}^{\alpha\alpha}(0)$ be real and positive can be successfully applied over the whole time domain to restore the correct mixing of real and imaginary parts. Cross-correlations are also typically well rephased with corrections computed from the corresponding autocorrelations. Two comments are needed here: on one hand, the fact that the phase correction can be extended at all times strongly suggests that the origin of the associated error on the physical hardware is of systematic nature. We will indeed argue in Sec. 2.3.3 that the phase inaccuracy is mostly due to systematic off-resonant driving of individual qubits resulting in biased rotations on the Bloch sphere. The validity and legitimacy of this rephasing procedure, and in particular its extension to more complicated cases where e.g. no exact benchmark solution is available, strongly relies upon such understanding of hardware errors. On the other hand, the manifestation of phase inaccuracies as apparent horizontal (i.e. temporal) shifts of the outcome dynamical correlations can easily be understood by noticing that for monochromatic oscillations of the form

$$f(t) = Ae^{i\omega t} = A(\cos \omega t + i \sin \omega t) \quad (2.39)$$

a constant phase shift can easily be interpreted as a time offset:

$$\hat{f}(t) = Ae^{i\omega t} e^{-i\phi} = A[\cos \omega(t - t_0) + i \sin \omega(t - t_0)] \quad (2.40)$$

where $\omega t_0 = \phi$. When, as it is often the case in our simulations, the oscillations result from different frequency components, the actual effect on the real and imaginary parts is more complicated.

The second step of the PaS procedure aims at mitigating the overall damping of the experimental data, mainly arising from incoherent errors on the quantum hardware. We make use of the following sum rule for $t = 0$ autocorrelations:

$$\langle (\hat{s}^{(i)})^2 \rangle = s^{(i)}(s^{(i)} + 1) = \sum_\alpha \mathcal{C}_{ii}^{\alpha\alpha}(0) \quad (2.41)$$

where $(\hat{s}^{(i)})^2$ is the square of the total i -th spin operator with value $s^{(i)} = 1/2$. A scaling factor F for the experimental data is then obtained as

$$F_i = \frac{s^{(i)}(s^{(i)} + 1)}{\sum_\alpha \mathcal{C}_{ii}^{\alpha\alpha}(0)} \quad (2.42)$$

This F_i can be used to magnify all autocorrelations $\mathcal{C}_{ii}^{\alpha\beta}(t)$. Moreover, we heuristically find an average between F_i and F_j to be a good scaling correction for $\mathcal{C}_{ij}^{\alpha\beta}$ cross-correlations. It is essential to notice that such scaling factors depend on the total length of the quantum circuit used to obtain the raw data, as longer circuits mean longer exposure of qubits to environmental noise: it must then be computed for each series of data depending e.g. on the underlying target Hamiltonian model and the number of Trotter steps. However, given that the duration of the single and two qubit gates

Table 2.3: Oscillation frequencies and most relevant Fourier coefficients for spin dimers obtained by fitting the PaS-mitigated data computed on the ibmqx4 chip. In square brackets, we report the corresponding exact results obtained from the numerical diagonalization of the spin Hamiltonian: the agreement is very good.

	Molecule 1	Molecule 2	Molecule 3
ω_1/J	2.00(2) [2.00]	9.5(1) [9.40]	9.5(1) [9.50]
ω_2/J	3.00(3) [3.00]	12.1(1) [12.10]	12.0(1) [12.00]
$ \langle 0 s_x^{(1)} 1\rangle ^2$	0.125(2) [0.125]	0.24(2) [0.24]	0.25(1) [0.25]
$ \langle 0 s_x^{(1)} 2\rangle ^2$	0.125(5) [0.125]	0.02(1) [0.01]	0.00(1) [0.00]
$ \langle 0 s_x^{(2)} 1\rangle ^2$	0.125(5) [0.125]	0.02(1) [0.01]	0.00(1) [0.00]
$ \langle 0 s_x^{(2)} 2\rangle ^2$	0.125(5) [0.125]	0.24(2) [0.24]	0.25(1) [0.25]
$\langle 0 s_x^{(1)} 1\rangle\langle 1 s_x^{(2)} 0\rangle$	-0.127(5) [-0.125]	-0.05(1) [-0.05]	0.00(1) [0.00]
$\langle 0 s_x^{(1)} 2\rangle\langle 2 s_x^{(2)} 0\rangle$	0.127(5) [0.125]	0.05(1) [0.05]	0.00(1) [0.00]

on the IBM Q processors is constant and does not depend on the continuous phase parameters, the scaling does not depend on the simulated time t , and can indeed be applied over the whole time domain to all data points sharing the same underlying structure of the quantum circuit. In Fig. 2.7c we show the combined effect of PaS corrections on the data for molecule 1 corresponding to the raw results in Fig. 2.7b. Both rephasing and rescaling parameters must always be computed for a consistent set of data taken in a particular moment and on a certain quantum processor: indeed, they intimately reflect, in addition to the characteristics of the particular sequence of gates that is being run, the calibration properties of the quantum hardware and the particular choice of the physical qubits. Moreover, it is worth pointing out explicitly that the data at $t = 0$ which are used as a reference point for the PaS procedure must be computed by setting the parameter t to zero in the corresponding quantum circuit but leaving all other quantum gates, for example CNOT appearing in the decomposition of $U_{ij}(t)$ blocks, untouched³.

After the experimental results have been treated with the PaS procedure, data are fitted to extract Fourier amplitudes

$$A_{ij,p}^{\alpha\beta} + iB_{ij,p}^{\alpha\beta} = \langle 0|s_\alpha^{(i)}|p\rangle\langle p|s_\beta^{(j)}|0\rangle \quad (2.43)$$

and frequencies E_p , see Eq. (2.30), which directly represent the most important physical information encoded in the quantum dynamical correlation functions under study, namely spin matrix elements and eigenenergies. In particular, we assume that each correlation function is a linear combination of oscillating functions

$$\begin{aligned} C_{ij}^{\alpha\beta}(t) &= \sum_p \left[A_{ij,p}^{\alpha\beta} + iB_{ij,p}^{\alpha\beta} \right] e^{-iE_p t} \\ &= \sum_p \left[A_{ij,p}^{\alpha\beta} \cos E_p t + B_{ij,p}^{\alpha\beta} \sin E_p t \right] \\ &\quad + i \sum_p \left[B_{ij,p}^{\alpha\beta} \cos E_p t - A_{ij,p}^{\alpha\beta} \sin E_p t \right] \end{aligned} \quad (2.44)$$

Frequencies and Fourier coefficients are extracted by combining a Fast Fourier Transform analysis with the FMINUIT package⁴. In the case of spin dimers, we always find that only two normal frequencies have non-negligible weight, while three frequencies are required to reproduce the measured oscillations. The finite time domain of the quantum simulation and the relatively small number of Trotter steps all limit in principle the resolution with which the exact frequencies can be obtained: however, having in mind a comparison with direct experimental measurements of dynamical correlations on real quantum systems, it is often sufficient that this resolution matches the experimental

³That is to say that $t = 0$ data points are *not* obtained by simply omitting the time evolution part in the quantum circuit: on the contrary, the $U_{ij}(0)$ block is used essentially as a noisy identity operation.

⁴See http://www.fis.unipr.it/~giuseppe.allodi/Fminuit/Fminuit_intro.html for details.




Figure 2.8: Major classes of hardware errors on IBM Q devices. (a) Incoherent Errors arise from uncontrolled interaction of the qubits with the environment. (b) Systematic coherent errors arise from non-perfect implementation of unitary gates. Here we show the effect of a single qubit off-resonant rotation (ORR): the ideal trajectory for a $R_x(\theta)$ gate (in red) is compared to an effective rotation around an axis tilted by $\theta_{\text{tilt}} = \pi/16$ towards the z direction (blue dots). The Bloch sphere plot is obtained with the Qutip python library [204, 205]. (c) Measurement errors induce erroneous qubit readout in the computational basis, with e.g. a classical outcome 1 when the qubit is actually in state $|0\rangle$.

one. In our case, we find excellent agreement with exact numerical Hamiltonian diagonalization, as it is reported in Tab. 2.3 for spin dimers and in Appendix A. In Fig. 2.7d-e we also show a selection of PaS mitigated results for molecule 2 and 4 with the corresponding fit. As mentioned above, for molecule 2 we used $n = 2$ Trotter step in the first part of the evolution ($Jt \leq 2.0$) and $n = 4$ afterwards: PaS mitigation was performed independently on the two data series, while the final fit was done over the whole time domain. More experimental data are reported for completeness in Appendix A.

2.3.3 Error analysis and scalability

Before proceeding to show how the information retrieved from dynamical correlation functions simulated on a quantum processor can become useful for further physical studies, we provide in this section a more detailed analysis of the hardware errors and of their relationship with the experimental results reported above. We stress that the following study was conducted from the perspective of IBM Q end users without any access to the physical hardware other than cloud connection. To compare and corroborate our findings, we also report numerical results obtained with the Qiskit simulation feature, which allows to run quantum circuits on classical hardware and under the effect of different noise channels⁵. The main error sources on IBM Q devices can be organized in three classes (see Fig. 2.8 for a pictorial representation):

- **Incoherent Errors (IEs)** are induced by the interaction with the external environment and are essentially due to non perfect isolation of the qubits. These are most often considered as the dominant limiting factors for quantum information processing applications, since they are essentially responsible for the finite lifetime (T_1) and coherence time (T_2) of quantum hardware, and are the target of many error-correction codes. In our simulations of hardware noise, incoherent errors are included by taking into account random qubit reset, applied with probability $p = e^{-t/T_1}$ after a noisy gate of duration t and mimicking an average $T_1 = 30 \mu\text{s}$, a Pauli-Z channel

$$\rho \mapsto p\sigma_z\rho\sigma_z + (1-p)\rho \quad (2.45)$$

⁵Latest Qiskit releases incorporate dedicated libraries for classical simulation (Qiskit Aer) and error analysis (Qiskit Ignis). The results presented in this section were obtained with an earlier version of the simulator, originally introduced in Qiskit v0.5.




Figure 2.9: Gate tomography for IBM Q processors. (a) Diagonal components of the quantum state $R_x(\theta)|0\rangle$ computed on ibmqx4 compared to ideal results (dashed lines). (b) Corresponding real and imaginary parts for off-diagonal components. These experimental data require pre-measurement U_{meas} unitaries to be obtained, as off-diagonal ρ elements are proportional to $\sigma_{x,y}$ expectation values while IBM Q devices only allow for measurements along the z direction: in this case, error propagation is taken into account to extract the typical tilt parameters (see text). (c) Experimental (ibmqx4, left column) and noisy numerical simulation (right column) of a CNOT gate applied to the initial quantum state $\sqrt{2}|\psi_0\rangle = |00\rangle + |10\rangle$. ORR, incoherent and measurement errors are included in the simulated data. (d) Same experiment as in panel (c), for initial state $\sqrt{2}|\psi'_0\rangle = |01\rangle + |11\rangle$.

to model pure dephasing [5] with $T_2 = 30 \mu\text{s}$ ⁶ and depolarizing Pauli channels

$$\rho \mapsto \frac{p_{\text{dep}}}{3} (\sigma_x \rho \sigma_x + \sigma_y \rho \sigma_y + \sigma_z \rho \sigma_z) + (1 - p_{\text{dep}}) \rho \quad (2.46)$$

applied with probability $p_U = 0.002$ and $p_{\text{CNOT}} = 0.05$ after each single- or two-qubit operation respectively. These values were based on average device calibration data and adjusted to match the experimental evidence. Here ρ indicates the density matrix

$$\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix} \quad (2.47)$$

of the numerically simulated system. Finite thermal population in the qubit excited state is found to be practically negligible, in agreement e.g. with Ref. [100].

- **Measurement Errors (MEs)** are the result of biased qubit readout, and typically cause a damping of the output quantities towards maximally mixed results. Qubit readout in superconducting setups, which is essentially based on detecting the frequency shift induced in

⁶Similarly to qubit excited state lifetime T_1 , the coherence time T_2 can be related to the probability p appearing in the definition of the Pauli-Z channel by assuming $p = e^{-t/T_2}$, where t is the duration of the noisy gate.

an auxiliary resonator via a microwave signal [206], is currently a limiting factor in many practical cases [166]. The measurement protocols are usually slow, with typical duration of several hundreds of nanoseconds, and their accuracy is on average around 96% [166], even though higher fidelities (above 99%) were reported in proof-of-principle demonstrations [207]. Asymmetric measurement errors are also possible, as in general the qubit excited state will be subject to relaxation while the readout is being performed. In our numerical simulations we assumed for simplicity symmetric measurement errors of 4-5%, in agreement with typical calibration data for IBM Q processors, see Tab. 2.1.

- **Systematic Coherent Errors (SCEs)** arise from non-perfect implementation of the coherent unitary dynamics corresponding to individual quantum gates. Concatenating several non-ideal quantum gates can significantly alter the dynamics under study. On IBM Q processors, we find the major manifestation of SCEs to be essentially consistent with Off-Resonant Rotations (ORR) of individual qubits, which have also been reported in the literature for similar devices [100]. ORR errors are typically due to inaccuracy in the frequency characterization of qubits, due to e.g. renormalization or drift of nominal values during operation cycles, and consequently in the design of driving pulses. We can examine an explicit example, adapted from Ref. [100], by looking at the trajectories on the Bloch sphere for an attempted $R_x(\theta)$ unitary gate reported in Fig. Fig. 2.8b. If the driving pulse is exactly on resonance with the qubit frequency Ω , the unitary time evolution of the qubit is

$$U(t) = e^{-i\frac{\Omega}{2}t\sigma_x} \quad (2.48)$$

and leads to the desired gate for $\Omega t = \theta$. If, on the contrary, the driving is detuned from Ω by an amount Δ , the actual time evolution is

$$\tilde{U}(t) = e^{-it(\frac{\Omega}{2}\sigma_x + \frac{\Delta}{2}\sigma_z)} \quad (2.49)$$

The rotation is now no longer performed around the x coordinate axis, but instead around a direction tilted towards the z axis by an angle $\theta_{\text{tilt}} = \tan^{-1}(\Delta/\Omega)$. Similarly, a wrong phase in the complex amplitude of the driving field can also lead to an effective rotation axis with an acquired component in the y direction. As it can be seen in Fig. Fig. 2.8b, the net effect is often a combination of unexpected offsets both in the x - y plane, leading to a wrong mixing of real and imaginary components of the off-diagonal entries in the qubit density matrix, and along the z direction, leading to artificial damping of oscillations. These effects are clearly related to the experimental errors reported in the quantum computation of dynamical correlations, see Sec. 2.3.2. To assess the impact of these ORR-induced inaccuracies on IBM Q devices, we performed a full $R_\alpha(\theta)$ tomography, whose typical results are reported in Fig. 2.9a-b for the case $\alpha = x$. Apart from an angle-independent measurement error of about 4-5%, which is explicitly manifest in the fact that e.g. $\rho_{00}(0) \neq 1$ when the initial state is $|0\rangle$, the ORR induces an effective damping⁷ of the oscillations and a wrong mixing of the Bloch vector components in the x - y plane. The latter correspond to $\sigma_{x,y}$ expectation values, which are proportional to the real and imaginary part of ρ_{01} entries. From these characterization experiments, we find typical effective rotation axis to be systematically tilted of approximately $\pi/16$ towards z and of $\pi/8$ ($-\pi/12$) towards y (respectively x) for all R_x (respectively R_y) gates. We recall that rotations around the z axis are operated virtually at software level and are thus error free by design.

Two-qubit SCEs can also in general arise from non-ideal realization of the native CNOT gate, for example due to spurious cross-talk effects. In our case, we performed a full gate tomography by preparing the initial states

$$\frac{|00\rangle \pm |10\rangle}{\sqrt{2}}; \quad \frac{|01\rangle \pm |11\rangle}{\sqrt{2}} \quad (2.50)$$

⁷Notice for example that $\rho_{11}(\pi) < \rho_{00}(0)$ while equality should hold.

and applying a CNOT operation, which ideally should bring the quantum register into the well know Bell basis. We then extracted, with suitable combined measurements⁸, all the components of the density matrix of the outcome state. This procedure requires several single-qubit rotations in addition to the two-qubit CNOT under examination, each one affected by its own SCEs: by comparing the results obtained in real IBM Q processors with numerical results including all the IEs, MEs and ORR errors described above, we find that single-qubit SCEs are already the dominant error source, as the noise model reproduces well the experimental findings, see Fig. 2.9c-d. Typical coherent errors for cross-resonance gates, described e.g. in Ref. [104] have a smaller impact and were thus neglected in subsequent numerical simulations.

To summarize, our remote diagnostic experiments on IBM Q processors provided us with a noise model featuring, as dominant error sources, single-qubit SCEs caused by off-resonant rotations. These constitute the major source of inaccuracy in the quantum simulation of dynamical correlation functions, as they can in principle completely alter the unitary dynamics on the quantum register and e.g. affect the information about the system eigenfrequencies and eigenstates. Other important error contributions come from incoherent errors and measurement errors, which usually manifest themselves as damping of the computed dynamical correlations. Classical numerical simulations including hardware errors are reported in Fig. 2.10. First, in Fig. 2.10a we demonstrate, by a direct comparison with experimental data for molecule 1, how phase and attenuation errors are very well reproduced by the theoretical model. As we discussed extensively, PaS mitigation can be used to partially recover the correct dynamics and, in the cases presented here, to extract correct fit parameters: however, as the net effect of SCEs usually goes beyond simple systematic phase shifts, increasing the complexity of the simulation often leads to residual alterations which are not fully corrected. In Fig. 2.10b-c we show some of such residual errors displayed in PaS-mitigated experimental data for molecule 4: by comparing with the results of noisy simulations, first performed with full IEs and SCEs contributions and then repeated with IEs only, it is straightforward to prove that SCEs are indeed the sources of these non-trivial dynamical inaccuracies. IEs alone are indeed mostly responsible of overall damping of the oscillations, which is easily fixed with the PaS procedure. Finally, we present in Fig. 2.10d experimental and simulated data for the case of a 4-spin Heisenberg model: the digital quantum simulation in this case involves 5 qubits (4 spins plus the ancilla) and was performed on the ibmqx4 quantum processor.

It is fair to say that, on the present IBM Q hardware, SCEs are currently the most important limitation for the digital quantum simulation of spin-spin dynamical correlation functions. We explicitly notice that, while quantum error correction (QEC) is of course the only viable approach to counterbalance qubit finite lifetime and decoherence, the errors we experienced in practical simulations are mostly an issue of qubit control and manipulation, and are not easily fixed by standard QEC techniques. On one hand, error mitigation strategies can certainly help in reducing a posteriori the impact of systematic errors, even though a fair and sensible application of these procedures requires at least a deep understanding of the origin and effect of the errors themselves. On the other hand, and contrary to intrinsic T_1 and T_2 limitations, SCEs could be at least partially tackled by quantum engineers with improved chip design, calibration and control electronics⁹. To assess the overall scalability of our method in a realistic scenario, we performed a series of numerical simulations targeting an open linear spin chain with isotropic nearest-neighbour $H_{\text{Heis},2}^{i,i+1} \equiv H_{\text{Heis},2}^i$ Heisenberg interactions

$$H_{\text{Heis},2}^i = J \sum_{\alpha=x,y,z} \sigma_{\alpha}^{(i)} \sigma_{\alpha}^{(i+1)} \quad (2.51)$$

⁸With combined measurements we mean here single-qubit unitary operations U_{meas} followed by qubit readout in the computational basis, which is the only available measurement process on IBM Q processors. As described in Sec. 1.1.5, this is sufficient to target e.g. the x or y components of the Bloch vector, i.e. the expectation values of $\sigma_{x,y}$. By measuring all possible combinations $\langle \sigma_{\alpha} \sigma_{\beta} \rangle$ a 2-qubit state can be fully reconstructed.

⁹We can for example imagine that a slightly sparser CNOT connectivity could lead to less inter-qubit spurious interactions, cross-talks and frequency renormalization corrections: empirical evidence in our quantum simulations suggests for example that smaller processors (ibmqx4) have better performances in terms of SCEs with respect to larger and more connected ones. This is indeed the route that IBM seems to be undertaking with the a new series of quantum processors, see e.g. the IBM Q Pughkeepsie 20-qubit chip in Chapter 4.




Figure 2.10: Error analysis and scalability. (a) Comparison between experimental raw data (top plot) and numerical simulations (bottom plot) featuring the estimated hardware errors (IEs, SCEs, MEs). Attenuation and phase errors are very well reproduced. (b) Experimental PaS-mitigated data for molecule 4. Residual errors not compensated by the PaS procedure are highlighted with blue circles. (c) Noisy simulations for the same dynamical correlation function of panel (b), after applying PaS corrections. The dynamics remains slightly altered if SCEs are present (top plot), but is fully recovered from error mitigation as soon as they are turned off (bottom plot). (d) Extension to the $N = 4$ isotropic Heisenberg chain. Experimental data are obtained on the ibmqx4 quantum processor. (e) Dynamical correlation functions for N -spin isotropic Heisenberg chains, simulated numerically including current levels of hardware noise (top), halved SCEs (middle) and IEs only (bottom). The quantum simulations require $N + 1$ qubits, including the ancilla, and are performed with $n = \text{round}[2\sqrt{N} - 2]$ digital Trotter steps.

examining, for increasing total system size N , the dynamical correlations on one end of the chain (i.e. for spins on the first and second site)¹⁰ Cross-correlation results are reported in Fig. 2.10e: in all simulations, the total gating time per digital step does not increase for $N > 3$, since pairwise interactions across even and odd bonds along the chain can be simulated in parallel as in the following example

$$\begin{array}{c}
 \boxed{e^{-iH_{\text{Heis},2}^{12}t/n}} \\
 \boxed{e^{-iH_{\text{Heis},2}^{34}t/n}} \\
 \dots \\
 \boxed{e^{-iH_{\text{Heis},2}^{45}t/n}}
 \end{array}
 \quad (2.52)$$

¹⁰This choice ensures that, even if the digital quantum simulation is performed in full for all the N spins, the observed quantity does not change too much with increasing N . Indeed, the effect of distant additional spins at the opposite end of the chain will only slightly affect the dynamical correlations on the first two sites, in such a way that we can focus only on the propagation of errors.

However, in order to keep a constant total digital error (and thus to reproduce the physical dynamics with sufficient accuracy) the number n of Trotter steps must increase with N . Indeed, by assuming a first order Trotter approximation and recalling Eq. (1.67), the total digital error is given by

$$\epsilon_N = \frac{\tau^2}{2} \left[\sum_i H_{\text{Heis},2}^{2i-1}, \sum_i H_{\text{Heis},2}^{2i} \right] \quad (2.53)$$

where $\tau = t/n$. As a result, noticing that all Heisenberg interaction terms have the same structure, ϵ_N scales as

$$\epsilon_N \simeq \frac{\tau^2 J^2}{2} (N-2) \left[\sum_{\alpha=x,y,z} \sigma_\alpha^{(1)} \sigma_\alpha^{(2)}, \sum_{\alpha=x,y,z} \sigma_\alpha^{(2)} \sigma_\alpha^{(3)} \right] \propto \frac{t^2}{n^2} (N-2) \quad (2.54)$$

In the numerical simulations, we fix ϵ_N to remain equivalent to the digital error arising in the quantum computation of the $N = 3$ cases (e.g. molecules 4 and 5) with $n = 2$, thus obtaining a slow increase

$$n = \text{round}[2\sqrt{N-2}] \quad (2.55)$$

of the number of digital steps with the size of the chain. With present error levels, the maximum N for which a reliable digital simulation could be obtained, after applying PaS mitigation and data smoothing through running average, is $N = 6$. Simply halving the magnitude of SCEs errors (i.e., in our case, the amplitude of θ_{tilt} angles) while keeping current hardware IEs could already extend the range up to $N \simeq 12$, and the results become of course very good after removing all SCEs. While these numbers are still relatively far from the threshold of pure quantum advantage¹¹, they look promising in terms of future scalability of the method as soon as larger and better quantum processors become available. It is also worth mentioning that some practical quantum speed-up for specific applications should be taken into account as a short- and mid-term goal. For example, we will show in the next section how the information retrieved from the digital quantum simulation of dynamical correlations can be used to simulate real experiments investigating the structure of magnetic molecular clusters: in everyday practice, most of the related data analysis tasks involve several repeated fits of Hamiltonian parameters, thus requiring on the numerical side to recompute many times the spin-spin dynamical correlations with slightly modified couplings and energies. Even if not impossible in principle, this is usually very demanding on classical devices already for moderately large system sizes: being able to carry out similar computations with sufficient accuracy on a quantum processor could offer the opportunity of reducing the required time and memory resources, enabling a better and more efficient use of experimental results and facilities. It goes without saying that for all those systems whose size exceeds the available classical computational resources a quantum simulation approach would be the ultimate solution, even though it is more difficult to foresee a realistic time scale for such achievements at present.

2.3.4 Four-dimensional inelastic neutron scattering

Molecular magnetic systems are ideal test-beds to investigate the phenomenology of finite-size spin systems. They typically contain several magnetic centers giving rise to rich and complex physical behaviors, determined e.g. by the spatial arrangement of the centers themselves along with interactions inside isotropic magnetic cores and anisotropic few-body perturbations [208]. Any attempt to model molecular magnets usually relies heavily on many of the paradigmatic spin Hamiltonians that we have thoroughly described and reviewed in the context of digital quantum simulation: however, the main difficulty in matching theoretical description and experimental practice in the study of single molecule magnets typically lies in determining, by comparing data and numerical simulations, the appropriate structural parameters, including e.g. the set of interaction energies and exchange parameters. Moreover, as the size of the target system increases, classical computational methods suffer from the exponential increase of the resources required by the corresponding quantum mechanical description. As a consequence, even though nowadays large magnetic molecules can be

¹¹Indeed, the numerical simulations reported here were all obtained on standard classical computers in a few hours.

chemically synthesized [209–212], a full understanding of the resulting properties and an interpretation of spectroscopic measurements is still lacking in many cases of interest. It is therefore natural to ask whether an approach based on quantum computing, and specifically relying on the near-term processors which will with greater probability dominate the scene in the next decade, could provide new resources in this challenge.

Along this chapter, we have presented a proof-of-principle demonstration of how NISQ prototype processors, in combination with classical data analysis and error mitigation techniques, can be used to digitally simulate dynamical correlation functions and extract the physical structural parameters of the corresponding spin lattice models. Here we show how these procedures can be turned into useful tools to tackle a practical problem of interest: in particular, we will use the data obtained for molecules 1-5 in Sec. 2.3.2 to compute cross sections for the so called four-dimensional inelastic neutron scattering experiments. These quantities are simultaneously typically hard to obtain on classical computers for large clusters and directly comparable with experimental evidence.

Four-Dimensional Inelastic Neutron Scattering (4D-INS) is the technique of choice to analyze molecular magnetic clusters [213–215], the reason for the name lying in the simultaneous dependence on both energy and three-dimensional momentum of the resulting spectra. By transforming back from reciprocal space, 4D-INS can therefore provide an unrivaled amount of information about the temporal and spatial dynamics of the normal excitations inside the target system. The cross section for 4D-INS experiments can be expressed, at temperature $T = 0$, in the general form [216]

$$I(E, \mathbf{Q}) \propto \sum_{i,j} \mathcal{F}_i(Q) \mathcal{F}_j^*(Q) \sum_{\substack{p \\ \alpha=x,y,z}} \left(\delta_{\alpha\beta} - \frac{Q_\alpha Q_\beta}{Q^2} \right) \langle 0 | s_\alpha^{(i)} | p \rangle \langle p | s_\beta^{(j)} | 0 \rangle e^{-i\mathbf{Q} \cdot \mathbf{R}_{ij}} \delta(E - E_p) \quad (2.56)$$

where E and \mathbf{Q} are the energy and three-dimensional momentum variables, $\mathcal{F}_i(Q)$ is the magnetic form factor for the i -th ion inside the molecule and \mathbf{R}_{ij} is the vector of relative positions of the i -th and j -th magnetic centers. While form factors and ion positions are usually known, eigenenergies E_p and spin matrix elements between eigenstates typically represent the computational bottleneck, as their calculation essentially requires the full diagonalization of the Hamiltonian model describing the magnetic molecule under study. However, by recalling Eq. (2.30), it is straightforward to observe that the Fourier frequencies and expansion coefficients of spin-spin dynamical correlation functions provide direct access to these quantities. We can then propose the following hybrid quantum-classical strategy to compute 4D-INS cross sections: first, a quantum processor efficiently performs the digital quantum simulation of all two-spin dynamical correlations $C_{ij}^{\alpha\beta}(t)$. Second, the results undergo classical post-processing for error mitigation via the PaS procedure, after which a classical Fourier transform or a fit analysis extracts the Fourier frequencies and coefficients. Finally, these are combined with known quantities to obtain the 4D-INS cross section as a function of transferred energy and momentum, which can then be compared with experimental characterization data for the target molecule. As mentioned at the end of the previous section, this calculations must usually be repeated several times, adjusting Hamiltonian parameters in order to match the experimental evidence from 4D-INS results, and quantum processors could either allow to simulate what is currently beyond reach of classical machines alone (currently limited in practice to a few dozens of spins) or provide useful speed-ups for smaller system sizes.

With the data obtained on IBM Q quantum processors for the digital quantum simulation of spin-spin correlation functions, presented in the previous sections and in Appendix A, we can now readily present a few examples of the proposed technique for 4D-INS cross section calculations. For both dimers (molecules 1-3) and trimers (molecules 4-5), we use realistic magnetic form factors and inter-atomic distances for Cu^{2+} ions. In particular, we assume the structure of dicopper diporphyrines [217] with $\mathbf{R} = (5 \text{ \AA})\mathbf{x}$ for dimers and equilateral triangles with 5 \AA edges for trimers. The results for the energy and momentum dependence of 4D-INS cross sections is presented in Fig. 2.11 for molecules 1-3 and Fig. 2.12 for molecules 4-5. It is worth mentioning explicitly that, as it is customarily done in 4D-INS experiments, all intensities $I(E, \mathbf{Q})$ are given relative to the maximum: this eliminates any dependence on the overall scaling factors and/or damping in the quantum simulation of dynamical correlations, as long as the structure and the oscillation patterns are respected. Moreover, in computing the final cross-sections reported here we made use of the




Figure 2.11: 4D-INS cross sections for spin dimers. (a) Molecule 1: The intensity $I(E)$ is reported as a function of energy (red dots in the top panel) and of momentum for $E = E_p$ of the two energy peaks (color maps in left column for peak I, right column for peak II), compared with exact results (solid line in the top panel, bottom color plots). The three dimensional plot for $I(E, \mathbf{Q})$ is also shown, while insets close to the energy peaks depict the corresponding spatial precession pattern of spin excitations. Similar results are reported in panel (b) for molecule 2 and in panel (c) for molecule 3.

fact, which was confirmed by experiments on the real quantum processors reported for completeness in Appendix A, that for the spin models under study $\mathcal{C}_{yz}^{\alpha\beta}(t)$ correlations are identically zero, $\mathcal{C}_{ij}^{xy}(t) = -\mathcal{C}_{ij}^{yx}(t)$ and $\mathcal{C}_{ij}^{zz}(t)$ are constant in time. Axial ($x \leftrightarrow y$) and permutational ($i \leftrightarrow j$) symmetries of the target Hamiltonian, when present, lead to the equivalences $\mathcal{C}_{ij}^{xx}(t) = \mathcal{C}_{ij}^{yy}(t)$ and $\mathcal{C}_{ij}^{\alpha\alpha}(t) = \mathcal{C}_{ji}^{\alpha\alpha}(t)$ respectively: whenever possible, we took advantage of this general properties by using, for example, $\mathcal{C}_{ij}^{xx}(t)$ instead of $\mathcal{C}_{ij}^{yy}(t)$. Indeed, the latter involves a larger number of gates to implement controlled- σ_y operations with respect to the hardware-native controlled- $\sigma_x = \text{CNOT}$, and is thus more error prone.

By comparing with ideal results computed numerically, it is immediately clear that the agreement is good in all cases. For spin dimers, the results in Fig. 2.11a-c show almost perfect resemblance with the exact counterparts, and we will show in the next section how the cross section data can be used to quantify two-spin entanglement. While the energy peaks identify the transitions to excited eigenstates, the extracted \mathbf{Q} -dependence of the inelastic neutron cross-section describes the structure of the eigenstates themselves. For instance, transitions between states belonging to different total spin multiplets are characterized by a minimum in $I(\mathbf{Q})$ for $\mathbf{Q} \rightarrow 0$. For example, in the case of molecule 1, Fig. 2.11a, the peak labeled as I corresponds to an inter-multiplet transition between the ground state triplet and the excited singlet $1/\sqrt{2}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$, while peak II represents a transition within triplet states, namely the ground $|\downarrow\downarrow\rangle$ and the excited $1/\sqrt{2}(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle)$. This is consistent with the positive sign of J in the target Hamiltonian, leading to a singlet state lower in energy than the triplet. Moreover, the \mathbf{Q} -dependence for fixed $E = E_p$ directly provides information about the spin precession pattern in real position space associated with the given excitation [215]. When a resonant perturbation brings a molecule from its ground state into a superposition state with a small amplitude on the p -th excited state, the spins described by the components ($\langle s_x^{(i)} \rangle, \langle s_y^{(i)} \rangle, \langle s_z^{(i)} \rangle$) start precessing around the z direction with frequency E_p . Looking again at molecule 1 and to the insets in Fig. 2.11a, the inter-multiplet transition (i.e. the first energy resonance) is characterized by the x and y components of the two spins precessing in phase opposition, while a giant spin excitation, with the two spins rigidly precessing conserving the same total-spin modulus of the ground state, is associated to the second energy resonance.

For molecules 4 and 5 the digital quantum simulation of dynamical correlations is more demanding and the estimated 4D-INS cross sections are slightly less accurate than in case of dimers. In particular, we notice that for molecule 4 (Fig. 2.12a) the relative intensities of the peaks in $I(E)$ suffer from slight deviations with respect to exact results. This is mainly due to the residual damping of \mathcal{C}_{21}^{xx} for $Jt > 8$, see Fig. 2.7e and Fig. 2.10b, which is not recovered from PaS-mitigation and leads to a Fourier coefficient $\langle 0|s_x^{(2)}|2\rangle\langle 2|s_x^{(1)}|0\rangle \neq 0$, as reported in Tab. A.1-A.2 (Appendix A). In real space, the corresponding excitation would induce a slight precession of ion 2, while exact diagonalization gives $\langle 0|s_2^x|2\rangle\langle 2|s_1^x|0\rangle = 0$ and ion 2 not precessing. For molecule 5, whose plots appear in Fig. 2.12b, we notice again slight deviations in the position of the first peak and in the relative intensities of the peaks in $I(E)$, while the $I(Q_x, Q_y)$ maps are in very good agreement with the exact behavior.

Finally, in addition to the scaling considerations expressed in Sec. 2.3.3, we point out that to compute the 4D-INS cross section only dynamical correlations between pairs of spins are needed. Accordingly, the total number of quantum gates and repetitions of the digital quantum simulations algorithm with different target qubits scales polynomially as N^2 with the total size of the system. The extension to molecular clusters containing magnetic centers with spin $s > 1/2$ can also be done by encoding the state of each spin s onto that of $2s$ qubits. For instance, four qubits $\sigma_{1,\dots,4}$ can represent a pair of interacting spins $s^{(1,2)} = 1$, i.e. $s_\alpha^{(1)} = (\sigma_\alpha^{(1)} + \sigma_\alpha^{(3)})/2$, $s_\alpha^{(2)} = (\sigma_\alpha^{(2)} + \sigma_\alpha^{(4)})/2$. In this particular example, the calculation of dynamical correlation functions for the target Hamiltonian would be recast in terms of correlation functions on pairs of physical qubits as

$$\begin{aligned} \langle s_\alpha^{(2)}(t)s_\beta^{(1)}(t) \rangle &= \frac{1}{4} \langle [\sigma_\alpha^{(2)}(t) + \sigma_\alpha^{(4)}(t)][\sigma_\beta^{(1)} + \sigma_\beta^{(3)}] \rangle \\ &= \frac{1}{4} \left[\langle \sigma_\alpha^{(2)}(t)\sigma_\beta^{(1)} \rangle + \langle \sigma_\alpha^{(4)}(t)\sigma_\beta^{(1)} \rangle + \langle \sigma_\alpha^{(2)}(t)\sigma_\beta^{(3)} \rangle + \langle \sigma_\alpha^{(4)}(t)\sigma_\beta^{(3)} \rangle \right] \end{aligned} \quad (2.57)$$

In this form, the desired quantities could then be directly evaluated on a quantum processor.




Figure 2.12: 4D-INS cross sections for spin trimers. (a) Molecule 4, digital quantum simulations performed on ibmqx4 and ibmqx5 processors. (b) Molecule 5, digital quantum simulations performed on the ibmq_20_tokyo processor.

2.3.5 Entanglement and concurrence

The information contained in 4D-INS spectra is not limited to the energetic and spatial structure of molecular excitations. For example, it has been shown that cross sections similar to the ones we described in the previous section can be used to portray entanglement in weakly coupled molecular qubits [214]. In our particular case, the \mathbf{Q} modulation of the scattered neutron intensity provides a quantification of entanglement in the excited states of a spin dimers characterized by a factorized reference ground state. In spin-1/2 dimers, i.e. for two-qubit states, entanglement is usually analyzed in terms of concurrence C [218–220], which is defined as follows. Let ρ be the density matrix of a (generally mixed) bipartite two-qubit state: after applying spin flip σ_y operations

$$\tilde{\rho} = (\sigma_y \otimes \sigma_y) \rho^* (\sigma_y \otimes \sigma_y) \quad (2.58)$$

the spectral decomposition of $\rho \tilde{\rho}$ is computed as

$$\rho \tilde{\rho} = \sum_{i=1}^4 \lambda_i^2 |\psi_i\rangle \langle \psi_i| \quad (2.59)$$

If the resulting eigenvalues are conventionally put in increasing order $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \lambda_4$, then the concurrence of ρ is defined as

$$C(\rho) = \max\{\lambda_1 - \lambda_2 - \lambda_3 - \lambda_4, 0\} \quad (2.60)$$

Maximally entangled states have $C(\rho) = 1$, while separable states correspond to $C(\rho) = 0$. In the particularly simple case of pure states $\rho = |\psi\rangle \langle \psi|$, where

$$|\psi\rangle = a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle \quad (2.61)$$

it is straightforward to compute an analytic expression for the concurrence

$$C(|\psi\rangle \langle \psi|) = 2|ad - bc| \quad (2.62)$$

At the same time, if we go back to molecules 1-3 and we concentrate on the first excited states (energy peak I in Fig. 2.11a-c, i.e. eigenenergy E_1 and eigenstate $|p\rangle = |1\rangle$), we can write

$$|\psi\rangle = |p\rangle = b|01\rangle + c|10\rangle, \quad C(|\psi\rangle \langle \psi|) = 2bc \quad (2.63)$$

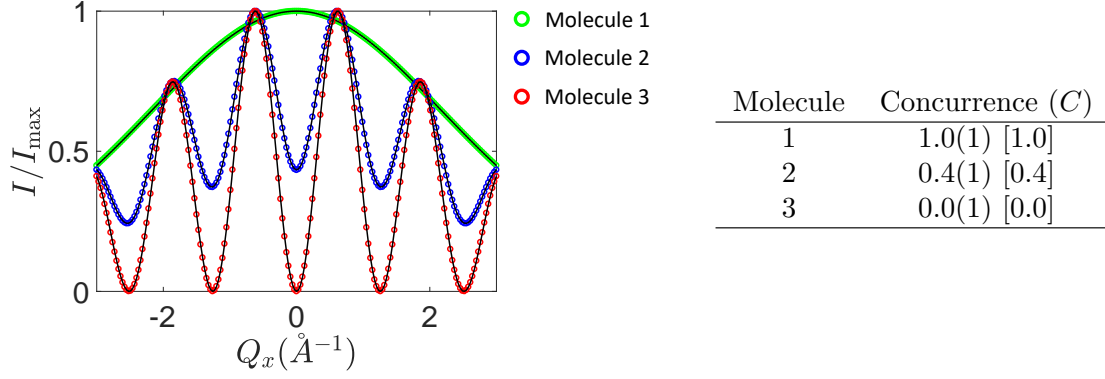


Figure 2.13: Concurrency fit for the first excited state in spin dimers. Cuts of the 4D-INS spectrum at $I(Q_x, Q_y = 0, Q_z = 0, E = E_1)$ are reported for molecules 1-3, where dots represent data obtained from IBM Q-based simulations and solid lines are the corresponding exact numerical results. Fitting the oscillations with Eq. (2.64) provides the concurrence values shown in the table on the right, which must be compared to the ideal values in square brackets.

and, recalling that we set $\mathbf{R}_{12} = R\mathbf{x}$,

$$\begin{aligned}
 I(Q_x, Q_y = 0, Q_z = 0, E = E_1) &\propto \mathcal{F}_1(Q_x)\mathcal{F}_2^*(Q_x)(|b|^2 + |c|^2 + bc^*e^{iQ_x R} + h.c.) \\
 &= \mathcal{F}_1(Q_x)\mathcal{F}_2^*(Q_x)[|b|^2 + |c|^2 + 2\text{Re}(bc^*)\cos(Q_x R) - 2\text{Im}(bc^*)\sin(Q_x R)]
 \end{aligned}
 \tag{2.64}$$

In the equations above we have already explicitly used the fact that correlations $\mathcal{C}_{ij}^{zz}(t)$ are independent of time to conclude that $d = 0$ ¹² and the presence of only two peaks in the energy spectra to exclude two-spin excitations (i.e. $a = 0$). A fit of the cross section \mathbf{Q} data (color maps in Fig. 2.11a-c, cut at $Q_y = 0$ and for peak I) then provides a direct estimate of parameters b and c and a measure of the concurrence, namely of the degree of entanglement, in the first molecular excited state. Results reported in Fig. 2.13 show an excellent agreement between the estimates based on IBM Q digital quantum simulations and the corresponding exact numerical calculations.

At a more fundamental level, quantum correlation effects can also be observed already from spin matrix elements: in particular, whenever the ground state $|0\rangle$ is separable a sizable cross-correlation coefficient $\langle 0|s_x^{(1)}|p\rangle\langle p|s_x^{(2)}|0\rangle$ fingerprints a genuinely quantum behavior due to entanglement in the excited state $|p\rangle$. Indeed, the values reported in Tab. 2.3 show that the cross-correlation coefficients decrease from molecule 1, which is characterized by maximum entanglement in both excited states, to molecule 3 in which excited states are factorized. Conversely, in molecules 2 and 3, autocorrelation coefficients $|\langle 0|s_x^{(i)}|p\rangle|^2$ are much larger for $(i, p) \in \{(1, 1), (2, 2)\}$ than for $(i, p) \in \{(1, 2), (2, 1)\}$, indicating that in these cases the two transitions are close to independent single-spin excitations. Finally, a much smoother, or even completely absent, modulation of $I(Q_x)$ in molecules 2 and 3 with respect to molecule 1 is also a signature associated to smaller entanglement, as it is clear from Fig. 2.13.

To conclude, in this chapter we have introduced an original combination of quantum algorithms, classical data analysis and error mitigation procedures through which we obtained reliable results for the digital quantum simulation of dynamical correlation functions. On one hand, this provides a clear experimental assessment of the current capabilities and limitations of cloud-accessible NISQ superconducting devices operated as universal quantum simulators. At the same time, we have

¹²Remember that we identified the factorized 2-spin ground state $|0\rangle = |\downarrow\downarrow\rangle$ with the qubit state $|11\rangle$ on the quantum register, and therefore a non-trivial time dependence in

$$\mathcal{C}_{ij}^{zz}(t) = \sum_p \langle 0|s_z^{(i)}|p\rangle\langle p|s_z^{(j)}|0\rangle e^{-iE_p t}$$

can only come from terms of the form $\langle 11|s_z^{(i)}|11\rangle$, which are in turn only present when $d \neq 0$.

also demonstrated a practical application in which the theoretical understanding of paradigmatic spin models and of the quantum protocols aimed at calculating their behavior in time is put into practice for the study of a concrete physical problem. With forthcoming technological progress, our approach could turn a non error-corrected quantum hardware into a valuable resource to speed-up the investigation of complex systems and unlock previously unexplored regimes, thus taking a further step into the NISQ-era roadmap towards quantum-enhanced technical computing.

Electromechanical Quantum Simulators

After exploring the near past and the present of quantum technologies for digital information processing and simulation, we will now take a look at possible future perspectives, with particular attention to hardware solutions. So far, we have seen that trapped ions and superconducting quantum circuits represent the dominant and most technically advanced platforms, though it is not clear whether any of the two will eventually represent the ultimate choice to realize a fully-fledged quantum computing architecture. It might also be the case that only a combination of several complementary approaches will lead to major advances. With this in mind, the first part of the chapter is dedicated to a brief review of alternative proposals and an outlook on the field of hybrid quantum technologies: we will then move forward and present a recent original theoretical study [92] combining electromechanical nano-oscillators and superconducting elements to provide an elementary building block for universal quantum simulations, with promising performances in terms of key figures of merit and scalability.

3.1 Perspective platforms for quantum information processing

Despite an impressive and fast development that brought superconducting qubits and trapped ions on the verge of a much hailed second quantum revolution, the road towards a functioning large scale quantum processor is still far from being clearly traced. On one hand, we have mentioned already in Chapter 1 how the extremely good coherence and connectivity of trapped ion qubits hide intrinsic limitations in scaling the linear chain design, while superconducting circuits are easily fabricated and manipulated but often suffer from a number of issues such as higher noise levels, gate infidelities, frequency crowding, cross talks, lack of connectivity and, last but not least, readout errors. Moreover, in practical realizations trapped ions require high and ultra-high vacuum conditions and superconducting qubits must be operated at cryogenic temperatures, thus limiting the versatility and portability of quantum computing devices. If the latter obstacles could easily be overcome, in perspective, with remote access¹, the physical and engineering challenges which separate us from the final goal of fault tolerant quantum computation will probably last several years or decades even in the most optimistic scenarios. In the meantime, intermediate-scale quantum devices represent both a necessary step and an opportunity for novel applications.

In parallel to the steady progress that the two leading technological platforms have been undergoing, a number of alternative approaches have been proposed over the past years, either running behind the scenes for some time after a few initial achievements or being studied *ex-novo*. The levels of maturity range from proof-of-principle demonstrations to more preliminary theoretical investigations. Taking into account the widespread scientific and commercial interest in practical quantum information processing, it would not be surprising to see in the near future some of these complementary solutions step out of the shadow of trapped ions and superconducting circuits and

¹This is indeed the vision that most quantum computing ventures and companies have already started to put into practice, from the IBM Quantum Experience to Rigetti's Quantum Cloud Services or D-Wave's Leap.

acquire an independent status as candidate platforms for quantum technological and computing purposes.

With their huge success in microelectronics and a number of appealing features, from room temperature operations to the potential for mass scale manufacturing and integration with existing classical technology, semiconductors have long been considered for practical quantum computing since the very beginning [221]. For example, semiconductor based quantum dots have been proposed and tested as spin qubits [222]. Single-qubit read-out and manipulation were demonstrated more than a decade ago [223, 224], though the intrinsic limitations to coherence times posed by nuclear spin dephasing and spin-orbit coupling [225] remain a severe challenge to scalability. Lately, promising experimental results in silicon-based quantum dots triggered a wave of renewed interest and moderate optimism: CNOT gates between two quantum dots have been shown with about 78% fidelity [226]. Precision for single qubit rotation can be as high as 99% [227–230], with average gating times below 100 ns and dephasing times close to 200 ns [228].

Controlled impurities and defect ions in silicon have also been considered for potentially low-cost quantum technologies [231]. As with quantum dot-based qubits, single-spin read-out and manipulation of localized donor impurities [232, 233] have already been achieved, and several proposals for two-qubit quantum gates have been put forward [234, 235]. Very recent results have demonstrated the first $\sqrt{\text{SWAP}}$ two-qubit operation between phosphorus donor electrons in silicon, with gating time of 800 ps and fidelity around 94% [236]. Along the same lines, color centers in diamond have been employed for quantum information manipulation and storage. As an example, Nitrogen-Vacancies (NV) were used to demonstrate decoherence-protected single- and two-qubit gates on a ten-qubit quantum register [237], achieving genuine multipartite entanglement over up to 7 spins. In the same work, NV centers were also proposed as stable multi-qubit memories, preserving single-qubit states up to 75 s and multi-qubit entanglement for over 10 s.

While photonic circuits have been largely explored as analog quantum simulators [70], it is less obvious how they could play a role in digital quantum computing and universal quantum simulations. Indeed, naturally weak photon-photon interactions, typically due to small material nonlinearities, complicate the construction of functioning quantum registers. Possible solutions might come from suitable electromagnetic confinement in nonlinear materials, which could lead to single-photon sensitivity [238, 239], or by using mixed radiation-matter excitations in semiconductors, also called exciton-polaritons, which proved to be sensitive at the single quantum level [240, 241]. Like conventional photonic circuits, exciton-polaritons have shown interesting applications in the field of analog simulators, particularly in the study of strongly interacting photonic lattices [19, 242]. In both photonic and polaritonic platforms, full on-chip integration could of course be a useful and necessary intermediate achievement towards quantum technological maturity. It is also worth mentioning that, within the now flourishing quantum computing commercial ecosystem, a few companies and start-up ventures such as Toronto-based Xanadu² are actively investigating photonic solutions for continuous-variable quantum computing [243–245], with special attention to perspective machine learning applications [246].

The long coherence times and high degree of chemical tunability make of magnetic molecules another promising potential platform for quantum information processing [47]. Manipulation is often achieved through electromagnetic pulses, through which it is for example possible to effectively control qubit-qubit interactions in suitably engineered structures of electronic [247] or nuclear [248] spins. The richness of the molecular Hilbert space could also be exploited to encode logical qubits, embedding quantum error correction within single molecules [249].

Finally, arrays of optically or magnetically trapped neutral atoms, typically excited in Rydberg states [46], represent a rising alternative for quantum simulations. Analog versions with more than 50 Rydberg atoms have already been reported [33], and programmable two-qubit entanglement has been shown with $\sim 97\%$ fidelity [250]. Although several challenges need to be addressed for general-purpose quantum computing applications, particularly related to the short coherence times as compared to trapped ions [251], an important advancement has been marked by the recent implementation, in parallel on several clusters of atoms in a one-dimensional array of optical

²See website at <https://www.xanadu.ai/>.

tweezers, of universal two- and three-qubit entangling gates on ^{87}Rb qubits. These were encoded in long-lived hyperfine ground states and coupled via excitation to strongly interacting Rydberg states, achieving controlled-phase gate fidelity above 97.4%.

3.2 Hybrid quantum technologies

The large number of quantum devices and different quantum technological platforms which have appeared over the years, showing varying strengths and suffering of sometimes complementary weaknesses, have eventually inspired the idea of experimenting with hybrid approaches [252, 253]. The philosophy supporting all these attempts, and particularly those aimed at developing resilient and general-purpose quantum computing systems, is indeed rather simple: it is very likely that, in the long run, a hybrid technology will be in the best position to simultaneously meet all the stringent requirements in terms of scalability (possibly in multi-dimensional arrays or distributed networks), chip-scale integration and high operational reliability.

An ideal infrastructure for quantum information processing should be able to carry out three major classes of tasks, namely storage, manipulation and transmission of quantum states, with all necessary inter-conversions between those stages, such as for example quantum memory readout. At the same time, it has become more and more clear that none of the currently available quantum technologies has the potential to execute all these tasks with sufficiently good performances and versatility. For example, photons are usually associated to communication protocols, electronic states in atoms or electric charges in semi- and superconductors can be employed to rapidly process information and weakly interacting electronic [254] or nuclear [255] spins may realize reliable quantum memories. Interfacing different kinds of elements, although very challenging in practice, holds promise for significant advantages. Loosely speaking, in most examples the coherent coupling between components, represented here by creation operators a^\dagger and b^\dagger , takes the general form [253]

$$H_{\text{coupling}} = g_{\text{eff}}(a^\dagger b + h.c.) \quad (3.1)$$

and the main difficulty often lies in achieving a sufficiently strong interaction coefficient g_{eff} . Spatial, energetic or impedance matching consideration play the dominant role in designing and optimizing hybrid quantum devices. In view of these issues, it is not surprising that among the first hybrid proposals we find the combination of superconducting qubits with long-lived quantum memories based on spin ensembles [256, 257], where the interaction between the two components is mediated by microwave cavity photons inside a common coplanar waveguide (CPW) resonator. Indeed, although the typical coupling strength³ between a single spin and a microwave photon is only around $g \simeq 2\pi \cdot 10 \text{ Hz}$, when the cavity field c is made to interact with a N -spin ensemble the relevant Hamiltonian is given by the Tavis-Cummings model

$$H_{\text{coupling}} = g\sqrt{N}(cS^\dagger + h.c.) \quad (3.2)$$

where S^\dagger represent collective spin excitations and the resulting effective coupling $g_{\text{eff}} = g\sqrt{N}$ for $N \simeq 10^{11}$ - 10^{12} can overcome the natural spin decoherence and cavity losses. Along the same lines, atomic and molecular ensembles have also been studied both as good quantum memories and optical interfaces with communication photons through e.g. Raman transitions [258–260]. Spin ensembles and photons in tunable resonators were recently proposed as the backbone of a novel architecture [110, 111, 261] with promising performances for universal quantum computing and digital quantum simulations [54, 262]. Finally, other examples of hybrid quantum technologies include Nitrogen Vacancy (NV) centers coupled to carbon nanotubes [263] and superconducting waveguides [264], and Rydberg atoms integrated on superconducting chips [265, 266].

3.2.1 Mechanical quantum devices

Particularly interesting within the hybrid quantum technological paradigm is the role played by mechanical degrees of freedom. In *quantum* opto- and electromechanics, collective oscillation

³Assuming $\hbar = 1$, we will use from now on frequency units to express couplings.

modes of micro- and nanoscale systems are investigated in the few-phonon regime, and have been proposed as promising alternatives to photonic devices for conversion and transmission of quantum information [267, 268], as support for quantum information manipulation [264, 269] and for sensing applications [270–272]. Most of such applications rely on the ability of mechanical interactions to create parametric photon-phonon interfaces whose coupling “bandwidth” ranges from optical to microwave frequencies. Indeed, the typical Hamiltonian describing both optomechanical radiation pressure interactions (featuring e.g. mobile mirrors in Fabry–Perot cavities) and electromechanical LC-frequency modulation (via e.g. oscillating capacitor membranes) has the form [253, 268]

$$H_{\text{mech}} = \omega_c a^\dagger a + \Omega_m b^\dagger b + g_0 a^\dagger a (b^\dagger + b) \quad (3.3)$$

where a is the electromagnetic field operator for the cavity resonator (either optical or electrical) and b is the annihilation operator for the mechanical mode. The structure of H_{mech} is readily derived from the bare Hamiltonian

$$H = \omega_c(x) a^\dagger a + \Omega_m b^\dagger b \quad (3.4)$$

if we assume that the mechanical displacement x influences the cavity resonance frequency. Indeed, at first order we can write

$$\omega_c(x) \simeq \omega_c + x \frac{\partial \omega_c}{\partial x} \quad (3.5)$$

and we recover Eq. (3.3) if we promote the position variable to a quantum mechanical operator in the usual way, i.e. $x \propto (b^\dagger + b)$, and we define

$$g_0 = x_0 \frac{\partial \omega_c}{\partial x} \quad (3.6)$$

where x_0 is the mechanical zero-point fluctuation amplitude. Other forms of the opto/electromechanical coupling are also known in special cases, depending e.g. quadratically on the displacement [273, 274] or involving qubit operators instead of cavity field, see Sec. 3.2.2. If the cavity is driven by an external laser at ω_L , the Hamiltonian above can be put in a rotating frame and linearized around the average cavity population $\bar{n} = |\bar{\alpha}|^2$ as follows

$$a = \bar{\alpha} + \delta a \quad (3.7)$$

giving

$$H' = -\Delta \delta a^\dagger \delta a + \Omega_m b^\dagger b + H'_{\text{int}} \quad (3.8)$$

where $\Delta = \omega_L - \omega_c$ and

$$H'_{\text{int}} \propto g_0 \sqrt{\bar{n}} (\delta a^\dagger + \delta a) (b^\dagger + b) \quad (3.9)$$

As a result, even if initially ω_c and Ω_m were far apart, the laser driving can be used to bridge the gap while also amplifying the bare interaction strength to $g_{\text{eff}} = g_0 \sqrt{\bar{n}}$. This property can make a single mechanical membrane very well suited as a coherent transducer between optical and microwave radiation [275] with potential applications in the quantum regime. Moreover, by choosing a red-detuned laser $\omega_L = \omega_c - \Omega_m$ and using the rotating wave approximation (RWA) in Eq. (3.8), it is straightforward to see that the dominant interaction term becomes

$$H''_{\text{int,RWA}} \simeq g_{\text{eff}} (\delta a^\dagger b + h.c.) \quad (3.10)$$

This is the analog of the quantum optical beam-splitter Hamiltonian and provides the basis for sideband cooling of the mechanical motion [276, 277]. On the other hand, the same reasoning in the blue-detuned case $\omega_L = \omega_c + \Omega_m$ leads to the equivalent of a two-mode squeezing Hamiltonian

$$H'''_{\text{int,RWA}} \simeq g_{\text{eff}} (\delta a^\dagger b^\dagger + h.c.) \quad (3.11)$$

which opens the way to parametric amplification and the creation of entanglement between photonic and mechanical degrees of freedom [278].

For the purposes of this chapter, we will leave aside the vast field of optomechanics to concentrate more on quantum electromechanical realizations. In the next section, we will briefly describe typical set-ups and a few experimental milestones in the field, such as the successful integration of mechanical degrees of freedom within superconducting quantum circuits, with the purpose of setting the stage for the subsequent description of an original architecture for quantum information processing featuring mechanical qubit encoding.




Figure 3.1: Electromechanical coupling in superconducting circuits. (a) Fundamental model of capacitive coupling between a mechanical oscillator and a LC resonator. A control transmission line is also explicitly sketched. (b) False color images of a realization of the electromechanical superconducting setup. Here the LC resonator is composed of a spiral inductor shunted by a parallel-plate capacitor, all made in Aluminium on a Sapphire substrate. The upper plate of the capacitor is suspended 50 nm above the lower plate and is free to vibrate. The mechanical oscillation frequency is $\Omega_m/2\pi = 10.69$ MHz, with a quality factor $Q_m \simeq 3.6 \cdot 10^5$. Images reproduced from Ref. [279].

3.2.2 Quantum electromechanics

Electromechanical interactions can arise in the microwave domain by capacitively coupling a mechanical element within a superconducting circuit. A paradigmatic model is shown in Fig. 3.1a, where one of the parallel plates of the capacitor in a lumped elements LC resonator is mechanically compliant [280]: the motion of the plate modifies the total capacity and thus modulates the characteristic frequency of the electrical circuit, in accordance to Eq. (3.5). A remarkable realization of this set-up with a drum-shaped suspended membrane capacitor is reported in Fig. 3.1b. If x is the mechanical displacement from the equilibrium position of the moving plate, d is the equilibrium separation between plates and C_0 is the corresponding equilibrium capacitance, we have in general the following form for the total capacity in the LC resonator [280]:

$$C_T(x) = C + C_0 \left(1 - \frac{x}{d}\right) \quad (3.12)$$

Here, we left room for a position-independent fraction C of the total capacitance. The total capacitive energy of the system is therefore

$$E_C = \frac{Q^2}{2C_T(x)} \quad (3.13)$$

where Q is the electric charge on the capacitor. If we expand to first order in the displacement we get

$$E_C \simeq \frac{Q^2}{2C_T(0)} + x \left(\frac{dE_C}{dx} \right)_{x=0} = \frac{Q^2}{2C_T} + \frac{\beta}{2dC_T} x Q^2 \quad (3.14)$$

where $C_T \equiv C_T(0)$ and $\beta = C_0/C_T$. We can thus write the total Hamiltonian for the circuit in Fig. 3.1a as [280]

$$H = \frac{p^2}{2m} + \frac{1}{2} m \Omega_m x^2 + \frac{\Phi^2}{2L} + \frac{Q^2}{2C_T} + \frac{\beta}{2dC_T} x Q^2 + \frac{1}{2} i(t) \Phi \quad (3.15)$$

Here, (p, x) are conjugate momentum and position coordinates for the mechanical mode with frequency Ω_m and effective mass m , (Q, Φ) are conjugate charge-flux variables for the circuit, L is the inductance and $i(t)$ is a time-dependent driving current. Canonical quantization of the electrical and mechanical oscillator variables⁴ $[x, p] = [Q, \Phi] = i$ leads to

$$H = \omega_c a^\dagger a + \Omega_m b^\dagger b + \frac{g_0}{2} (b + b^\dagger)(a + a^\dagger)^2 - i(\mathcal{E} e^{-i\omega_d t} + \mathcal{E}^* e^{i\omega_d t})(a - a^\dagger) \quad (3.16)$$

⁴As in all other parts of this work, we consider $\hbar = 1$.

where

$$\begin{aligned} a &= \sqrt{\frac{\omega_c L}{2}} Q + \frac{i}{\sqrt{2\omega_c L}} \Phi \\ b &= \sqrt{\frac{m\Omega_m}{2}} x + \frac{i}{\sqrt{2m\Omega_m}} p \end{aligned} \quad (3.17)$$

The cavity resonance is $\omega_c = 1/\sqrt{LC_T}$, while ω_d is the frequency of the drive and the coupling constant reads

$$g_0 = \frac{\beta\omega_c x_0}{2d} \quad (3.18)$$

Here, $x_0 = 1/(2m\omega_c)$ is the usual zero-point quantum fluctuation length for the mechanical mode. Finally, by going to a rotating reference frame via the unitary transformation

$$\mathbf{H}' = \mathbf{U}\mathbf{H}\mathbf{U}^\dagger - i\mathbf{U}\frac{d\mathbf{U}^\dagger}{dt} \quad \text{with } \mathbf{U} = e^{i\omega_d a^\dagger a t} \quad (3.19)$$

and making the rotating wave approximation we get

$$\mathbf{H}' = (\omega_c - \omega_d)a^\dagger a + \Omega_m b^\dagger b + g_0(b + b^\dagger)a^\dagger a + (\mathcal{E}'a^\dagger + h.c.) \quad (3.20)$$

namely the typical radiation pressure-like behavior as in Eq. (3.3). Typical experimental parameters, reported e.g. in Ref. [279], are $\omega_c/2\pi = 7.5$ GHz, $\Omega_m/2\pi = 10.69$ MHz, $x_0 = 4.1$ fm and $g_0/2\pi = 230$ Hz. Successful integration of mechanical devices with superconducting coplanar resonators led over the past years to a number of milestone achievements, from strong coupling [279] to ground-state sideband cooling of the mechanical motion [281], coherent state transfer [282] and entanglement with propagating electric signals [278].

In parallel to linear LC resonators, a second fundamental and necessary step in hybrid quantum electromechanics is the integration of mechanical modes with non-linear superconducting circuits, i.e. qubits. Milestone achievements were reported already in 2009 by LaHaye *et al.* [283], who performed spectroscopic measurements of a nanomechanical suspended bridge coupled to a microwave-driven Cooper-pair box (CPB) qubit, and in 2010 by O'Connell *et al.* [284], who obtained ground state cooling of a high-frequency piezoelectric bulk acoustic resonator inside a dilution fridge and the subsequent superconducting qubit-mediated single phonon control. Particularly interesting for our purposes is the full hybrid cavity QED set-up proposed e.g. by Pirkkalainen and collaborators [285, 286] for Ω_m in the MHz range and by Rouxinol *et al.* [287] for ultra-high mechanical frequencies $\Omega_m/2\pi \simeq 3.4$ GHz, in which a transmon qubit is simultaneously coupled to a microwave cavity and to a phonon mode in a micromechanical resonator. In general, a qubit-mechanical oscillator capacitive interaction is described by the following Hamiltonian [280]

$$\mathbf{H} = \Omega_m b^\dagger b + \frac{\epsilon}{2}\sigma_z + \frac{\delta}{2}\sigma_x + \lambda(b + b^\dagger)\sigma_z \quad (3.21)$$

where the qubit energy parameters δ and ϵ can be independently controlled by a voltage gate bias and a magnetic flux bias. Interaction strengths in the order of a few MHz [285] up to $\simeq 100$ MHz [286, 287] have been reported, while typical qubit energy scales are in the 1 – 10 GHz range. Notice that the electromechanical interaction term is not diagonal with respect to the bare qubit eigenstates: indeed, after a rotation into a new set of qubit variables (i.e. by diagonalizing the qubit sector and redefining the Pauli axes)

$$\begin{pmatrix} \tilde{\sigma}_x \\ \tilde{\sigma}_z \end{pmatrix} = \begin{pmatrix} \epsilon/\Omega_q & -\delta/\Omega_q \\ \delta/\Omega_q & \epsilon/\Omega_q \end{pmatrix} \begin{pmatrix} \sigma_x \\ \sigma_z \end{pmatrix} \quad (3.22)$$

where $\Omega_q = \sqrt{\epsilon^2 + \delta^2}$, we obtain a new Hamiltonian

$$\mathbf{H} = \Omega_m b^\dagger b + \frac{\Omega_q}{2}\tilde{\sigma}_z + g_{m,z}(b + b^\dagger)\tilde{\sigma}_z + g_{m,x}(b + b^\dagger)\tilde{\sigma}_x \quad (3.23)$$

It is now clear that, in the actual qubit eigenbasis, the mechanical coupling in general has both diagonal ($g_{m,z}$) and transverse ($g_{m,x}$) components. However, in many cases of interests it can be shown that $g_{m,z} \ll g_{m,x}$ [285]: we will then assume from now on that the relevant physical behavior is well captured by the familiar Jaynes-Cummings Hamiltonian⁵

$$H = \Omega_m b^\dagger b + \frac{\Omega_q}{2} \sigma_z + g(b + b^\dagger) \sigma_x \quad (3.24)$$

As we will later explore the dispersive regime $\Omega_m \ll \Omega_q$, we keep for completeness both rotating and counter-rotating terms without making explicit use of RWA-like approximations.

Several interesting classes of mechanical systems have been put forward for circuit QED applications: in Fig. 3.2 we show, together with the already mentioned Aluminium membranes [279, 285] and suspended Silicon Nitride [283] or Aluminium [287] structures, examples of carbon nanotubes (CNTs) [289–291] and graphene sheets [288, 292, 293]. Remarkably high quality factors ($Q_m = \Omega_m/\Gamma_m$ where Γ_m is the relevant damping rate) and coherence properties were reported for many of these devices, particularly for $\Omega_m/2\pi$ equal to a few tens of MHz. Moreover, dynamical tuning of the mechanical characteristic frequency over a range of 5-10 MHz has been demonstrated thorough external gate voltages for SiC beams [294], Al membranes [295], CNTs [289] and graphene structures [288, 292].

3.3 Electromechanical qubits for digital quantum simulations

The remarkable quality factors and the progress in fabrication and control of nanomechanical devices led to the idea of encoding and process quantum information onto nanomechanical qubits. The project is fascinating in many respects, not least because of the fact that the very origin of programmable computing machines can be traced back to the fully mechanical analytical engine proposed in the XIX century by C. Babbage⁶. An optomechanical set-up made of several nanoresonators coupled through a common high-finesse optical cavity was theoretically described by Rips and Hartmann in 2013 [269], reporting promising gate fidelities. Here we propose to bring this approach to a new stage in the electromechanical domain, leveraging the seamless integration with existing superconducting technology and the possibility of controlling electromechanical devices at the single phonon level. We envision a hybrid architecture to efficiently implement a digital quantum simulator, based on electromechanical elements coupled to superconducting circuits and designed in a modular fashion, thus featuring built-in perspectives for scalability.

In our set-up, anharmonic and tunable nanomechanical resonators (NRs) will play the role of digital quantum information carriers (i.e. qubits), while virtual fluctuations of superconducting non-linear resonators such as transmons will be employed to mediate two-qubit gates: as we will demonstrate with extensive numerical simulations performed with state-of-the-art parameters, this scheme puts into practical use the good relaxation (T_1) and coherence (T_2) properties of NRs, while the corresponding time scales for the superconducting elements become essentially irrelevant for information processing purposes. In the subsequent sections, we will argue that this electromechanical qubit encoding would not only allow achieving high gate fidelities above 99.9% but also, and most importantly, a ratio $T_2/T_{\text{gate}} > 10^4$, where T_{gate} represents the average single- and two-qubit overall gating time. The latter is a key figure of merit in view of the realization of digital quantum simulations, in which long sequences of concatenated operations are most often required.

3.3.1 Hybrid electromechanical building block

In Fig. 3.3 we report a schematic representation of the elementary building block of the proposed hybrid architecture: this fundamental unit is given by a pair of electromechanical NRs mutually coupled to a nonlinear circuit element, here assumed to be a transmon. A superconducting resonator cavity or a transmission line is also schematically represented and can be taken into account as a

⁵For simplicity, we drop the tilde signs from the qubit operators in the diagonalized Hamiltonian.

⁶See for example C. Babbage (ed. by M. Campbell-Kelly), *Passages from the life of a philosopher*, Rutgers University Press, New Brunswick, N.J., USA (1994).




Figure 3.2: Electromechanical devices. (a) Suspended silicon nitride nanostructure with $\Omega_m/2\pi = 58$ MHz and $Q_m \simeq 3 \cdot 10^4$. Reproduced from Ref. [283]. (b) Suspended aluminium membrane coupled to a transmon qubit (shown in green), $\Omega_m/2\pi = 71.842$ MHz, $Q_m \simeq 5500$. Image reproduced from Ref. [285] (c) Doubly clamped aluminium membrane, with $\Omega_m/2\pi = 88$ – 89 MHz, $Q_m \simeq 10^3$ (measured at 77 K, estimated $Q_m \simeq 10^5$ at milli-kelvin temperatures). Images courtesy of Prof. M. D. LaHaye, Syracuse University. (d) Multilayer graphene sheet, $\Omega_m/2\pi = 36.23$ MHz, $Q_m \simeq 1,59 \cdot 10^5$. Reproduced from Ref. [288]. (e) Suspended carbon nanotube (CNT), $\Omega_m/2\pi \simeq 300$ MHz, $Q_m \simeq 6,250$. Image reproduced from Ref. [289]. Other experiments with similar structures reported Q_m up to $5 \cdot 10^6$ [290].

further element to be weakly coupled to each NR for ground state cooling [281], which in our case becomes equivalent to qubit initialization, and readout. Contrary to the much more common usage as qubit, the transmon is present here just as an interaction mediator, and in perspective could also be employed for read out of each NR qubit state [284, 285]. In the following calculations and characterizations we will not explicitly include the transmission line in the model, but we point out already at this stage that a sequential repetition of the fundamental NR-transmon-NR unit with one or more common resonator cavities suggests a natural design for scaling up the architecture.




Figure 3.3: Scalable electromechanical building block for digital quantum simulations. (a) Schematic representation of the fundamental unit for quantum information processing: two electromechanical oscillators (NR) are mutually coupled through a transmon within a superconducting circuit. The NRs are controlled and tuned individually through external voltage biases, while external magnetic fields can be used to tune the transmon transition frequency. A superconducting transmission line/cavity resonator for NRs initialization and/or readout is also sketched. A common transmission line can also act as the backbone of a sequential architecture where the NR-transmon-NR block is repeated to scale up the quantum processor. (b) A possible equivalent circuit representation of the building block, with the transmission line schematically represented by the L_r and C_r elements.

One elementary unit of the electromechanical platform can be modeled with the free Hamiltonian

$$H_0 = \sum_i \left[\omega_i b_i^\dagger b_i + H_{\text{nl},i} \right] + \frac{\Omega}{2} \sigma_z, \quad (3.25)$$

where b_i (b_i^\dagger) represent bosonic annihilation (creation) operators for the mechanical resonators, see e.g. Eq. (3.17), and the transmon is described as a pure quantum two-level system [156] with frequency Ω . In order to isolate the two lowest energy levels in the otherwise harmonic spectrum of the nanoresonators we include a non-linear term $H_{\text{nl},i}$: this will make it possible to unambiguously define a qubit computational basis, as it is customarily done in weakly anharmonic superconducting circuits. In order to keep our analysis simple, we will only require a shift of the lowest lying Fock energy levels, i.e. a diagonal term on the Fock basis

$$H_{\text{nl},i} = U b_i^\dagger b_i^\dagger b_i b_i \quad (3.26)$$

We will discuss the possible ways of engineering such non-linear contribution in Sec. 3.3.4. Finally, the interaction between mechanical oscillators and the transmon is modeled, following Eq. (3.24), as

$$H_{\text{int}} = \sum_{i=1}^2 g_i (b_i + b_i^\dagger) \sigma_x \quad (3.27)$$

We will demonstrate that with a realistic choice of U , g_i and $\Delta_i = \omega_i - \Omega$, the dynamics is effectively restricted to the computational basis constituted by the $n_i = 0, 1$ Fock states of the nanoresonators, provided that they are initialized within the same manifold.

Our set-up is designed to work in the strongly dispersive regime: in the following, the electromechanical resonators frequencies will be set below 100 MHz, while the transmon frequency in the 2-10 GHz range. On one hand, we motivate this choice by referring to the experimental evidence reporting remarkably better coherence times for electromechanical resonators working in the few tens of MHz, see Sec. 3.2.2. On the other hand, as already mentioned above, this will leave the transmon essentially in its ground state during all information processing operations performed in low-occupancy bosonic states, with NR-NR interactions mediated by its virtual excitations. Notice that in principle such an energy mismatch does not allow to directly employ the rotating wave

approximation in H_{int} : we will thus keep for completeness all rotating and counter-rotating terms in our perturbative analysis, although we will later be able to demonstrate that the latter only minimally affect the qualitative behaviour of the system in the regime of parameters under study. Dissipation and pure dephasing effects will be fully included in our simulations via the master equation [296]

$$\frac{d\rho}{dt} = i[\rho, H_0 + H_{\text{int}} + H_{\text{control}}(t)] + \mathcal{L}_{TR}[\rho] + \sum_i \mathcal{L}_i[\rho], \quad (3.28)$$

where $H_{\text{control}}(t)$ introduces an explicit time dependence for dynamical tuning and driving of the different components, which will be better specified in the following sections. Lindblad terms act individually on the electromechanical NRs

$$\mathcal{L}_i[\rho] = \gamma_i \mathcal{D}(b_i)[\rho] + \gamma_{i,d} \mathcal{D}(b_i^\dagger b_i)[\rho] \quad (3.29)$$

and on the transmon

$$\mathcal{L}_{TR}[\rho] = \gamma_{TR} \mathcal{D}(\sigma_-)[\rho] + \gamma_{TR,d} \mathcal{D}(\sigma_z)[\rho] \quad (3.30)$$

respectively. Here ρ is the joint density matrix of the full system and $\mathcal{D}(O)[\rho] = O\rho O^\dagger - 0.5\{O^\dagger O, \rho\}$, with $\{A, B\}$ denoting the anticommutator $AB + BA$.

3.3.2 Effective qubit-qubit interactions

By assuming that the single-phonon non-linearity U is sufficiently strong (typically of a few MHz in our simulations) to allow the clear spectral isolation of a computational basis and that the system is operated deeply in the dispersive regime ($\Omega \gg \omega_i$, $|\Delta_i| \gg g_i$), useful insights about the behavior of the electromechanical qubits within a single building block can be obtained by constructing an effective perturbative Hamiltonian. Indeed, when the above conditions are met, we can consider H_{int} as a small perturbation to H_0 . Moreover, as long as we refer to low-occupancy Fock states, two clearly separated manifolds can be identified in the unperturbed spectrum of H_0 : a low-energy one containing total eigenstates of the form⁷ $|n_1 \downarrow n_2\rangle$, and a high-energy one in which the transmon is in the excited state, i.e. the set of $|n_1 \uparrow n_2\rangle$ eigenstates. The magnitude of the gap between the two manifolds is set by the transmon frequency Ω and is therefore of the order of a few GHz, much larger than the typical phonon energy. We can then safely apply second order perturbation theory and then restrict ourselves to the low-energy manifold, tracing over the transmon degrees of freedom while keeping the influence of virtual transmon excitations as effective interactions between the nanoresonators. The resulting effective Hamiltonian matrix elements can be obtained from the following identity

$$\begin{aligned} \langle \mu\nu | \tilde{H}_{\text{eff}} | \mu'\nu' \rangle &= \frac{1}{2} \sum_{m_1, m_2} \langle n_1 \downarrow n_2 | H_{\text{int}} | m_1 \uparrow m_2 \rangle \langle m_1 \uparrow m_2 | H_{\text{int}} | n'_1 \downarrow n'_2 \rangle \\ &\times \left[\frac{1}{\omega_1(n_1 - m_1) + \omega_2(n_2 - m_2) - \Omega} + \frac{1}{\omega_1(n'_1 - m_1) + \omega_2(n'_2 - m_2) - \Omega} \right] \end{aligned} \quad (3.31)$$

where the sum runs explicitly over all states of the full Hilbert space in which the transmon is in the excited state. With some abuse of notation, we understand that \tilde{H}_{eff} acts on a smaller Hilbert space in which we only keep the NRs degrees of freedom in the low-energy manifold, namely $|\mu\nu\rangle \mapsto |n_1 \downarrow n_2\rangle$ with $\mu \equiv n_1$ and $\nu \equiv n_2$. If we further restrict to the computational basis, i.e. $\mu, \nu = 0, 1$, while using bosonic states up to $m_i = 2$ in Eq. (3.31), the resulting effective Hamiltonian can be expressed in terms of Pauli matrices as

$$\tilde{H}_{\text{eff}} = \sum_{i=1}^2 \left(\frac{\lambda_i}{2} \sigma_z^{(i)} \right) + \frac{\Gamma}{4} \sigma_x^{(1)} \sigma_x^{(2)} + \text{const.}, \quad (3.32)$$

⁷Here we use $|n_1 \sigma n_2\rangle$ to denote the unperturbed eigenstate of the system in which the i -th nanoresonator is in the n_i -th Fock state and the transmon, which is treated here as a pure two-level system, is in its ground ($\sigma = \downarrow$) or excited ($\sigma = \uparrow$) state.

where

$$\Gamma = \frac{4g_1g_2\Omega(\omega_1^2 + \omega_2^2 - 2\Omega^2)}{(\Omega^2 - \omega_1^2)(\Omega^2 - \omega_2^2)} \quad (3.33)$$

is an effective NR-NR coupling, while

$$\lambda_i = \frac{-2g_i^2(\Omega^2 + \omega_i(2U + \Omega))}{(2U + \Omega + \omega_i)(\Omega^2 - \omega_i^2)} \quad (3.34)$$

are single-qubit energy renormalization shifts due to the dispersive interaction with the transmon. The effective transmon-mediated two-qubit interaction in Eq. (3.32) can actually be decomposed in two terms:

$$\frac{\Gamma}{4}\sigma_x^{(1)}\sigma_x^{(2)} = \frac{\Gamma}{8}\left(\sigma_x^{(1)}\sigma_x^{(2)} + \sigma_y^{(1)}\sigma_y^{(2)}\right) + \frac{\Gamma}{8}\left(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)}\right) \quad (3.35)$$

The first XX + YY interaction on the right hand side can straightforwardly be rewritten as an exchange Hamiltonian

$$H_{\text{ex}} \propto (\sigma_+ \otimes \sigma_- + \sigma_- \otimes \sigma_+) \quad (3.36)$$

coupling the $|10\rangle$ and $|01\rangle$ elements of the computational basis. As we briefly mentioned in Sec. 1.1.3, this kind of interaction is common to many superconducting architectures for quantum information processing and here will constitute the basis to build a universal set of quantum gates. The second XX – YY contribution is actually generated by the counter-rotating contributions in the original interaction Hamiltonian, and couples the $|00\rangle$ and $|11\rangle$ components. Given the large energy gap $G \simeq \omega_1 + \omega_2$ between them, the net effect of this term is negligible as long as the couplings g_i are not too strong, in such a way that $\Gamma \ll G$. It is nevertheless worth keeping in mind that deviations from an ideal excitation exchange behavior can arise under different conditions. For most practical purposes, we can thus rewrite the effective Hamiltonian as

$$H_{\text{eff}} = \sum_{i=1}^2 \left(\frac{\lambda_i}{2} \sigma_z^{(i)} \right) + \frac{\Gamma}{8} \left(\sigma_x^{(1)}\sigma_x^{(2)} + \sigma_y^{(1)}\sigma_y^{(2)} \right) + \text{const.} \quad (3.37)$$

In the following, we will use the H_{eff} derived above as a qualitative reference to predict and understand the behavior of the real system when employed as a minimal two-qubit register. However, in order to fully account for the actual complexity of the set-up, we will use the full model in Eq. (3.28) to perform all numerical simulations and quantitative analysis. This will avoid artificial biases with respect e.g. to the possible population leakages from computational basis during qubit manipulations, while also checking *a posteriori* the validity of the approximations leading to H_{eff} .

3.3.3 Single and two-qubit gates

One of the key ingredients to perform single and two-qubit gates is the dynamical tuning of ω_1 and ω_2 . This can be achieved by using external static and modulated electric fields, i.e. electrostatic potential energies V , acting locally on a single NR. The working principles of these control techniques have already been studied and demonstrated experimentally on a large variety of systems [288, 289, 292, 294, 297]. A systematic voltage bias can immediately be used to cancel out the effective frequency renormalization contributions λ_i in Eq. (3.37). For the rest of the discussion, we will thus assume, unless explicitly stated and without loss of generality, that the reference frequencies of the NRs remain the ω_i defined at the beginning: in the numerical implementations, this is achieved by introducing additional terms $-\lambda_i b_i^\dagger b_i$ in the model Hamiltonian.

As a first step towards the description of gate operations, we must identify an idle configuration for the building block: this is readily found if we assume that the two NRs are significantly detuned, i.e. $|\omega_1 - \omega_2| \gg \Gamma$. Under such conditions, the transmon-mediated interaction term in Eq. (3.37) is practically ineffective: hence, the two qubits are essentially decoupled and independent rotations of each of them can be implemented. Here, the use of a high-frequency transmon helps improving the two qubits decoupling.

Following similar implementations [269], single qubit R_z rotations can be performed by shifting the NRs oscillation frequency for the amount of time required to add the desired phase to the

$n_i = 1$ component of the wave-function. For example, a change of the qubit oscillation frequency by an amount $\delta\omega$ for a time interval δt can be implemented thorough a step-like pulse, the temporal switching of the external fields being only limited by the response time of the control electronics (typically in the ns timescale). The resulting

$$H_{\text{control}}(t) = H_i^z(t) = \delta\omega\Theta_{t_0}(\delta t)b_i^\dagger b_i \quad (3.38)$$

will induce a time evolution

$$U_{\text{control}}(t) = \exp\left(-i \int dt H_{\text{control}}(t)\right) \quad (3.39)$$

producing a phase

$$\phi_z = \int dt \delta\omega\Theta_{t_0}(\delta t) = -\delta\omega\delta t \quad (3.40)$$

on the $|1\rangle$ state. Here $\Theta_{t_0}(\delta t)$ is a unitary square function starting at t_0 with duration δt and the usual form of R_z operators is immediately recovered up to a global phase by observing that, on the computational basis, $b^\dagger b \equiv (1/2)(\mathbb{I} + \sigma_z)$. Other single-qubit rotations are obtained with an oscillating transverse field keeping a definite phase relationship with the quantum mechanical evolution of the system. By choosing

$$H_{\text{control}}(t) = H_i^{xy}(t) = V^{xy}(t)(b_i + b_i^\dagger) \quad (3.41)$$

with

$$V^{xy}(t) = \Theta_{t_0}(\delta t)V_0^{xy} \cos(\omega_i t + \theta) \quad (3.42)$$

one can achieve either R_x (for $\theta = 0$) or R_y (for $\theta = \pi/2$) rotations: indeed, if we restrict $H_{\text{control}}(t)$ on the $n_i = 0, 1$ computational basis and we express it in terms of Pauli matrices we have

$$H_{\text{control}}(t) = \Theta_{t_0}(\delta t)V_0^{xy} \cos(\omega_i t + \theta)(\sigma_-^{(i)} + \sigma_+^{(i)}) \quad (3.43)$$

Moving to the interaction picture with respect to $\omega_i b_i^\dagger b_i$, using the well known identity

$$\cos(\omega_i t + \theta) = \frac{e^{i(\omega_i t + \theta)} + e^{-i(\omega_i t + \theta)}}{2} \quad (3.44)$$

and dropping off-resonant terms we finally get

$$H_{\text{control}}^I(t) \propto \cos(\theta)\sigma_x^{(i)} + \sin(\theta)\sigma_y^{(i)} \quad (3.45)$$

The total rotation angle equals the area under the pulse modulating the oscillation

$$\phi_{xy} = \int dt V_0^{xy} \Theta_{t_0}(\delta t) \quad (3.46)$$

In practical realizations of $R_x(\alpha)$ or $R_y(\alpha)$ gates, the pulse shape can more generally be written as

$$V^{xy}(t) = A(t, t_0, \sigma)V_0^{xy} \cos(\omega_i t + \theta). \quad (3.47)$$

where V_0^{xy} denotes the amplitude scale of the pulse, while $A(t, t_0, \sigma)$ is a time-dependent modulation of the oscillatory part that describes the on/off switching of the gate. If $\zeta(t)$ is the Heaviside step function, one straightforward choice of this pulse envelope is of course a square pulse

$$A(t, t_0, \sigma) = \zeta(\sigma/2 - |t - t_0 - \sigma/2|) \equiv \Theta_{t_0}(\sigma) \quad (3.48)$$

starting at t_0 and lasting for $\sigma = \alpha/V_0^{xy}$, as used in Eq. (3.42). However, as we are not dealing with real two-level systems but rather trying to restrict the dynamics of a nonlinear harmonic oscillator to its ground and first excited levels, a more careful design of the pulse profile can lead to significant improvement of the performances. Indeed, the frequency spectrum of a cosine-like

function modulated by a square pulse may in general not be sufficiently narrow around the target ω_1 to avoid the activation of unwanted transitions that are close in energy, e.g. the $|1\rangle \leftrightarrow |2\rangle$ transition. In this context, it is well known that an elementary but already powerful optimization tool is provided by Gaussian pulses, leading to a frequency spectrum which is also gaussian and rapidly decaying. The envelope can then be chosen as

$$A(t, t_0, \sigma) = e^{-\frac{(t-t_0)^2}{2\sigma^2}} \quad (3.49)$$

with $\sigma = \alpha/(\sqrt{2\pi}V_0^{xy})$. The gate will now last approximately $2.5-3\sigma$ on both sides of the central peak at $t = t_0$. In this way, the amount of nonlinearity that is required to obtain reasonably large gate fidelities can be reduced with respect to the square pulse case, or alternatively the strength of the tuning V_0^{xy} can be increased, thus diminishing the total gating time.

To realize two-qubit gates, the transmon-mediated interaction between the electromechanical NRs must be activated. On one hand, good isolation of the qubits during the idle phase and single-qubit gates requires that they are detuned from each other and that the transmon is at a sufficiently high frequency (e.g. $\Omega \simeq 10$ GHz) to strongly suppress the residual effective coherent coupling. On the other hand, dynamically tuning the two qubits to resonance, for example to a common intermediate frequency $\omega_r = |\omega_1 - \omega_2|/2$ is already sufficient to activate coherent oscillations: in this case, the time evolution induced by the relevant term in H_{eff} reads

$$U_{XY}(t) = e^{-i\frac{\Gamma}{8}(\sigma_x^{(1)}\sigma_x^{(2)} + \sigma_y^{(1)}\sigma_y^{(2)})t} \quad (3.50)$$

which, expressed in matrix form on the computational basis states becomes

$$U_{XY}(t) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \frac{\Gamma t}{4} & -i \sin \frac{\Gamma t}{4} & 0 \\ 0 & -i \sin \frac{\Gamma t}{4} & \cos \frac{\Gamma t}{4} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (3.51)$$

This is an entangling operation which, if the time evolution is halted at $\tau = \pi/|\Gamma|$, leads to the universal two-qubit gate

$$\sqrt{i\text{SWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1/\sqrt{2} & i/\sqrt{2} & 0 \\ 0 & i/\sqrt{2} & 1/\sqrt{2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (3.52)$$

The interaction is then turned off by shifting back the frequencies of the nanoresonators. The protocol is completed by two rephasing single qubit z -rotations applied to the NRs to correct for the additional phase accumulated by the qubits with respect to their idle evolution with the original bare $\omega_{1,2}$. This latter step is obtained by an inverse $-\xi_i = \omega_i - \omega_r$ pulse lasting for a time $\tau' = \text{mod}(\tau, 2\pi)$. The typical gating time of a single $\sqrt{i\text{SWAP}}$ is of the order of a few μs with realistic choices of the parameters. This duration can be optimized by exploiting one peculiar property of our setup, namely the possibility of dynamically tuning the transition frequency Ω of the transmon. This is, of course, of great importance when performing long sequences in the presence of realistic dissipation processes. This shift, which could be implemented experimentally by varying the magnetic flux concatenated with the transmon, is essentially just another time-dependent contribution to the total Hamiltonian in the interaction picture

$$H_{\text{TR}}(t) = \delta\Omega(t)\frac{\sigma_z}{2} \quad (3.53)$$

Changing the frequency of the transmon affects all the effective qubit-qubit Hamiltonian parameters: in particular, reducing Ω (and thus the detuning Δ with respect to the nanoresonators) increases the coupling $\Gamma \simeq g^2/\Delta$ and modifies the renormalizations λ_i . Needless to say, this procedure is limited both by the tunability range of the transmon and by the validity of the perturbative expansion in terms of g^2/Δ . Given some values of the external dissipation rates, for example, there exists an




Figure 3.4: Examples of single- and two-qubit quantum gates. (a)-(b) $R_x(\pi)$ rotation performed on qubit 1 while qubit 2 is kept isolated and the transmon frequency $\Omega = 10$ GHz is left unchanged, with the corresponding Gaussian oscillating pulse. The peak amplitude is 0.3 MHz. (c)-(d) $\sqrt{i\text{SWAP}}$ operation, with a short idle phase before and after the gate. Frequency shifts shown in (c) are operated on qubit 2 (NR pulse) and on the transmon (TR pulse) during the time evolution. A short rephasing stage on the qubit is included at the end, while the renormalization shift λ_2 is not explicitly reported. Similar operations are also performed on qubit 1 to achieve the resonance condition. In all panels, the system undergoes unitary evolution with bare NR frequencies $\omega_1 = 85$ MHz and $\omega_2 = 75$ MHz. In (b)-(d), $p_{ab} = |\langle ab|\psi\rangle|^2$ represents occupation probabilities for the Fock states computed on the simulated wavefunction $|\psi\rangle$ after tracing out the transmon degrees of freedom.

optimal $\delta\Omega$ that increases the gate fidelity without losing too much of the agreement between the actual behavior of the system and what is expected from the effective Hamiltonian description. In our simulations, we set a non zero $\delta\Omega$ during the coherent interaction time τ , which must then be computed by using $\Omega + \delta\Omega$ in Eq. (3.33) and we put the transmon back to its original frequency already during the qubit rephasing stage. We mention explicitly that when the transmon frequency is modified, the permanent shifts $-\lambda_i$ applied to the qubits must in principle be adjusted accordingly.

The set made of single qubit rotations and $\sqrt{i\text{SWAP}}$ is universal for standard quantum computation. In addition to that, the actual parametric nature of the native two-qubit operation will prove extremely useful for digital quantum simulation protocols, where the ability to activate the $U_{XY}(t)$ interaction for an arbitrary duration will avoid some of the more complicated constructions which are typically needed when using fixed-phase two-qubit gates. Notice that the set $\{R_\alpha, U_{XY}\}$ is exactly the \mathcal{S}_2 introduced in Chapter 1, see Eq. (1.38). A few representative examples of simulated single and two-qubit gates, together with a description of the required external pulses, are shown in Fig. 3.4. We notice that square pulses modulating NRs and transmon frequencies are already sufficient to obtain satisfactory results with two-qubit operations, as the U_{XY} evolution

conserves the total number of excitations and is thus less prone to population leakages outside the computational basis. Instead, gaussian envelope functions are usually required to obtain good single qubit rotations. With these precautions, we could verify that in all numerical simulations bosonic occupancies $n_i > 1$ do not significantly appear throughout typical gate dynamics.

The performances of the set-up were assessed in terms of the fidelity [5]

$$\mathcal{F} = \sqrt{\langle \psi | \rho | \psi \rangle} \quad (3.54)$$

between the ideal target pure state of a given single- or two-qubit operation (ψ) and the actual density matrix ρ resulting from the numerical integration of the full master equation, see Eq. (3.28). On one hand, this measures the resilience of our protocols with respect to dissipation and decoherence both on the NRs and on the interposed transmon. At the same time, the overall fidelity is also in principle affected by any residual unitary evolution effects not fully captured by the perturbative analysis which guided the design of the pulse sequences: possible examples are the non perfect isolation of the NRs in the idle phase, leading to spurious cross-talk effects, or excitation of the bosonic states with more than one phonon.

As benchmark operations, we report in Fig. 3.5 the results obtained for a $R_x(\pi/2)$ -rotation on qubit 1 and for the \sqrt{i} SWAP gate as a function of the pure dephasing rate (i.e. the reciprocal of the coherence time) for both the NRs and the transmon. For these simulations, the the system was initialized in a random superposition of all the 4 states in the computational basis, and we chose the following set of parameters: $\omega_1/2\pi = 85$ MHz, $\omega_2/2\pi = 75$ MHz (idle configuration), $U/2\pi = 3$ MHz, $g_1 = g_2 = 2\pi \cdot 6$ MHz, $\gamma_1 = \gamma_2 = 2\pi \cdot 50$ Hz, $\gamma_{TR}/2\pi = 100$ kHz. The transmon frequency $\Omega/2\pi$ is kept at 10 GHz in idle configuration and during single-qubit operations, and is tuned down to 2.5 GHz when performing two-qubit gates. As a direct consequence of the fact that the relevant quantum information content is carried by the nanoresonators, while the transmon is involved only through virtual excitations, the fidelity \mathcal{F} shows very weak dependence on the transmon decoherence rate $\gamma_{TR,d}$ and a much more sensitive dependence on $\gamma_{NR,d}$. Indeed, the simulations suggest that our scheme is in principle intrinsically robust against transmon decoherence, as the results look practically insensitive, for both single- and two-qubit gates, to an increase of more than two orders of magnitude in $\gamma_{TR,d}$ above the most optimistic [161] value $\gamma_{TR,d}/2\pi = 10$ kHz, corresponding to a transmon T_2 time of 100 μ s. Notice also that the typical gating time of a single \sqrt{i} SWAP is of the order of a few μ s, which is comparable or even larger than the shortest transmon T_2 time considered, i.e. $T_{2,TR} = 0.1$ μ s when $\gamma_{TR,d}/2\pi = 10$ MHz. Finally it is worth pointing out that a value of $\gamma_{NR,d}/2\pi \simeq 100$ kHz looks rather pessimistic for most candidate electromechanical qubits, in particular nanomembranes and nanotubes, where total linewidths in the 0.1-1 kHz range have been experimentally shown [288, 290, 292].

To compare with existing technology, we recall that on state-of-the-art transmon based technology the shortest two-qubit gating times are currently around $T_{\text{gate}} = 40$ -50 ns when fidelities above 99% are required [60], leading to a typical ratio $T_2/T_{\text{gate}} \sim 10^3$ with the current order of magnitude for superconducting qubit coherence times [161]. In our simulations, we obtained $\mathcal{F} > 99\%$ with $T_{\text{gate}} = 500$ (300) ns for single- (two-) qubit gates, thus demonstrating that the proposed electromechanical platform could potentially achieve $T_2/T_{\text{gate}} > 10^4$ with a still realistic coherence time around $T_{2,NR} \sim 10$ ms.

Residual thermal occupancy of nanomechanical resonators

Throughout our analysis, we assumed that the system can be cooled at sufficiently low temperatures to safely neglect the effect of thermal interaction with the environment. While conventional dilution refrigerators can easily maintain thermal noise below typical transmon excitation energies, achieving the corresponding regime for nanoresonators with fundamental frequencies in the MHz range is still challenging from the experimental point of view. Indeed, such relatively low-frequency devices most often require special techniques such as sideband cooling to reach their mechanical ground state [281]. For this reason, we performed an additional study on the gate fidelities considering, on top of all the intrinsic dissipation mechanisms reported in Eq. (3.28), a residual thermal interaction between the nanoresonators and the surrounding environment. In this series of simulations,




Figure 3.5: Fidelity of single- and two-qubit quantum gates. (a) Single qubit $\pi/2$ x -rotation. (b) Two-qubit \sqrt{i} SWAP gate. The fidelities are evaluated as functions of the pure dephasing rates $\gamma_{1,d} = \gamma_{2,d} = \gamma_{NR,d}$ of the electromechanical resonators and of the transmon ($\gamma_{TR,d}$), starting from a random superposition of all computational basis states. All other parameters are reported in the text.

we added to the master equation additional Lindblad terms describing a bosonic thermal reservoir attached to the i -th oscillator

$$\mathcal{L}_i[\rho] = \frac{\chi}{2} [\bar{n}(\omega_i, T) + 1] (2b_i\rho b_i^\dagger - b_i^\dagger b_i\rho - \rho b_i^\dagger b_i) + \frac{\chi}{2} \bar{n}(\omega_i, T) (2b_i^\dagger\rho b_i - b_i b_i^\dagger\rho - \rho b_i b_i^\dagger) \quad (3.55)$$

where ($k_B = 1$)

$$\bar{n}(\omega_i, T) = \frac{1}{\exp(\frac{\omega_i}{T}) - 1} \quad (3.56)$$

is the Bose occupancy at temperature T . In particular, we set $\bar{n} = 0.1$ for both oscillators, as this generally represents a good estimate of the achievable residual thermal occupation in real experiments. The rate χ can be inferred from typical line broadening of the nanoresonators, which is of the order of few tens of Hz. When average values for all the other dissipation mechanism are taken into account, the impact of the residual thermal occupation is of the order of 0.01-0.1% on single gate fidelities if $\chi = 50$ Hz, and around 1% if we use the more pessimistic value $\chi = 1$ kHz.

3.3.4 Non-linearity of electromechanical qubits

A sufficiently strong non-linearity at the single-phonon level is of crucial importance for the definition of electromechanical qubits, namely for a faithful encoding of digital quantum information in a spectrally isolated computational basis. In our previous analysis, such degree of anharmonicity was introduced through a simplified model featuring, in the Fock number representation, a diagonal shift of the $|1\rangle \leftrightarrow |2\rangle$ transition with respect to $|0\rangle \leftrightarrow |1\rangle$. While this treatment already contains all the features relevant for quantum computing purposes, keeping at the same time the description easy to understand and analytically transparent, it says little about how the non-linearity is actually obtained in practice. In this section, after assessing the actual impact of non-linearity on the gate fidelity, we will discuss the issue more extensively, first comparing our simplified description with a more realistic and commonly used model of anharmonic mechanical oscillators, and then by introducing an original hybrid solution to engineer the spectral properties of electromechanical resonators.

Fidelity versus non-linear shift

The relevance of the spectral anharmonicity for the fundamental quantum information processing tasks carried out by our proposed electromechanical set-up is best investigated with respect to single




Figure 3.6: Effect of non-linearity. (a) Fidelity of a single qubit $R_x(\pi)$ gate as a function of the nonlinear shift between the first and second energy gaps. Gaussian tuning pulses were used to activate the gate. (b) Corresponding population leak outside the computational basis. We plot the norm of $(\mathbb{I} - P_{\text{cb}})|\psi\rangle$, where $|\psi\rangle$ is the state of the system after the rotation (tracing out the transmon degrees of freedom) and P_{cb} is the projector on the computational basis subspace.

qubit transverse (i.e. x, y) rotations, as they require a form of driving which in principle does not conserve the total number of bosonic excitations in the system. As an example, in Fig. 3.6, we show how the fidelity of a single qubit $R_x(\pi)$ rotation changes as a function of the non-linear gap difference $\delta = \omega_{21} - \omega_{10}$, where $\omega_{10} = \omega_1 - \omega_0$ is the distance in frequency between the computational basis eigenstates ($|0\rangle, |1\rangle$) and $\omega_{21} = \omega_2 - \omega_1$ is the separation of the first non-computational excited state from $|1\rangle$. Notice that, within our simple non-linear model of Eq. (3.26), we have $\delta = 2U$. The simulations were carried out varying the parameter U in the absence of external dissipations. The most significant feature is the large plateau very close to unity for almost all values above $\delta \simeq 2$ MHz. Moreover, we find that a nonlinear shift $\delta = 1$ MHz is already sufficient to give theoretical fidelities $\mathcal{F} > 0.99$, which is the threshold required on single gates to successfully apply error correction protocols such as the surface code [60].

Models of anharmonicity

From the perspective of the total bosonic Hilbert space (i.e. without truncation on the maximum number $n_{1,\text{max}}$ and $n_{2,\text{max}}$ of allowed excitations in each resonator), the diagonal non-linearity model which we have adopted above can be written, for a single oscillator, as a Kerr-type Hamiltonian

$$H_{\text{nl,diag}} = Ub^\dagger b^\dagger bb. \quad (3.57)$$

On the other hand, a widely accepted model for nonlinear nanomechanical resonators is rather given by [298]

$$H_{\text{nl,full}} = U(b^\dagger + b)^4 \propto \hat{x}^4. \quad (3.58)$$

As a first step in comparing the two models, we will now show that, in the range of interest, the first one overestimates the magnitude of the parameter U required to achieve a given amount of non-linearity in the spectrum, namely it predicts, given U , a smaller shift $\delta = \omega_{21} - \omega_{10}$, with respect to Eq. (3.58). In Fig. 3.7, the shift δ for the two models is computed by performing a numerical diagonalization of the Hamiltonian, after adding the free term $H_0 = \omega b^\dagger b$ and setting $n_{\text{max}} = 10$. In the plot, U and δ are both expressed as fractions of the bare frequency of the oscillator ω . When $U/\omega \leq 0.1$, which roughly corresponds to the assumptions made in our previous description of electromechanical qubits, the model Hamiltonian (3.57) is always conservative in terms of quantitative estimation of the nonlinear single-phonon contribution given U . When the intrinsic mechanical non-linearities [294, 297] are not sufficiently strong, the anharmonic regime of nanoresonators can be accessed by stretching the structure out of its natural equilibrium position




Figure 3.7: Comparison between models for single-phonon nonlinearity of nanomechanical resonators. (a) Spectral shift $\delta = \omega_{21} - \omega_{10}$ predicted using a simplified diagonal model (solid line) and using a full non-linear Hamiltonian (dashed red line) for the same choice of the parameter U controlling single-phonon non linearity. (b) Focus on the region of interest for the typical set of parameters used in this work, namely $U/\omega \leq 0.1$.

via static external electric fields [294, 298–300]: the above results then suggest that the actual amplitude of such external bias required to obtain e.g. the target $\delta = 1$ MHz (see Fig. 3.6) might be experimentally accessible in real devices which obey the full non-linearity model reported in Eq. (3.58), even though single-phonon sensitivity has not yet been fully demonstrated. We will come back to this later when discussing alternative techniques to engineer a non-linear electromechanical spectrum.

For completeness, we also compared the eigenvectors corresponding to the first three energy levels (namely the computational basis plus the first extra excited level) for the two models $H_{\text{nl,diag}}$ and $H_{\text{nl,full}}$. In this case, we selected the parameter U in two different ways such that the gap shift δ is the same both for $H_{\text{nl,diag}}$ and $H_{\text{nl,full}}$. The following table summarizes the fidelity $\mathcal{F} = |\langle n_{\text{nl}} | n \rangle|$ of the eigenvectors $|n_{\text{nl}}\rangle$, obtained from Eq. (3.58) and ordered by increasing energy, with respect to the corresponding bare Fock state $|n\rangle$, for the two different choices $\omega = \omega_1 = 85$ MHz and $\omega = \omega_2 = 75$ MHz:

n	$\mathcal{F}(\omega = \omega_1)$	$\mathcal{F}(\omega = \omega_2)$
1	0.9996	0.9995
2	0.9971	0.9963
3	0.9899	0.9872

If the computational basis is redefined on the set of $|n_{\text{nl}}\rangle$ states, the very same machinery that we presented in Sec. 3.3.3 can be applied to implement single-qubit rotations and the \sqrt{i} SWAP gate. Indeed, electrical pulses can still be used to tune the fundamental transition frequency, ω_{01} , for both oscillators, thus bringing them to resonance when needed. Moreover, the operators b_i and b_i^\dagger still promote transitions between the new eigenvectors, albeit with a slightly different matrix element $X_{kl} = \langle k | b | l \rangle$, as summarized in the following table:

X_{kl}	Fock states	$ n_{\text{nl}}\rangle$ for $(\omega = \omega_1)$	$ n_{\text{nl}}\rangle$ for $(\omega = \omega_2)$
X_{01}	1	1.0005	1.0007
X_{12}	$\sqrt{2}$	1.4170	1.4179

The only effective consequence for the gating protocols is thus the (in principle transition-specific) rescaling of the coupling elements $g_i \rightarrow g_i X_{xy}$, leading to a renormalization of the required gating time or pulse amplitude. Numerical simulations give overall fidelities that are within 0.1% with respect to the simplified case using Fock states and $H_{\text{nl,diag}}$. It is also worth noticing that the

Hamiltonian in Eq. (3.58) only couples Fock states differing by an even number of excitations. This in turn means that $|n_{nl}\rangle$ and $|(n+1)_{nl}\rangle$ are still orthogonal to each other as in the bare Fock basis: indeed, they can be regarded as eigenvectors of the number parity operator belonging to different eigenspaces. In simpler terms, they are superpositions containing, alternatively, either even or odd Fock states. As a direct consequence, the matrix element X_{nn} vanishes. The same is not true for $X_{n(n+2)}$, meaning that b_i and b_i^\dagger could in principle promote $|0\rangle_{nl} \leftrightarrow |2\rangle_{nl}$ transitions outside the computational basis, for example during the single qubit x or y rotations. However, in our protocol this effect remains negligible in view of the large energy gap between the two distant eigenstates.

Inducing phonon non-linearity with a superconducting circuit

In Fig. 3.6, we saw that the degree of anharmonicity which would be required to achieve a reliable qubit behavior in our set-up goes from 1 to a few MHz, and we have already mentioned a few examples in the literature where the non-linear regime is reached via electric field gradients, see Refs. [294, 298–300]. Despite a few early investigations relying on those techniques for quantum information applications [269], it is still unclear if such approaches will ultimately be sufficient to generate reliable sensitivity of the non-linear effects down to the single-phonon level. We hereby explore an alternative scheme based on dispersive coupling of an additional low-frequency superconducting element to each NRs. We will demonstrate how this weak hybridization of quantum electromechanical excitations could in principle engineer by design a non-linear spectrum without degrading the remarkable coherence properties which constitute one of the main predicted strengths of nanomechanical qubits.

A combined system made of a single harmonic nanoresonator and a superconducting (SC) non-linear element would in principle display anharmonic energy levels which could be used for the definition of the physical qubit. Indeed, this coupled NR-SC can again be effectively described by a Rabi-like interaction term

$$H = \omega_{NR} b^\dagger b + \frac{\Omega_{SC}}{2} \sigma_z + g(b + b^\dagger)(\sigma_- + \sigma_+) \quad (3.59)$$

As done previously and in similar works [301], no rotating wave approximation is done at this stage, in view of selecting a dispersive regime $\omega_{NR} \ll \Omega_{SC}$. We stress that the SC component envisioned here would by no means play the role of a mediator between different NRs, as the transmon introduced in the fundamental building block of our set-up was doing (see Sec. 3.3.1 and Fig. 3.3). Instead, it would only be an additional element attached to each individual electromechanical oscillator with the explicit purpose of slightly modifying the structure of its excitation spectrum. This difference in roles manifests itself in the parameters range required to obtain a significant non-linear effect, which are very different from the ones at which a transmon qubit is typically operated. In fact, a SC resonance frequency Ω_{SC} in the few hundreds of MHz range is required here, as well as a NR-SC element coupling rate g in the few tens of MHz, which is basically opposite to the large Ω (1-10 GHz) and small g (a few MHz) previously assumed for the transmon mediator. From the experimental point of view, a suitable non-linear SC fluxonium circuit with low frequency and rather good dissipation and coherence properties compared to typical SC performances has already been demonstrated [302]. At the same time, quite large NR-SC coupling strengths can be obtained capacitively [285–287], for example by applying a sufficiently strong voltage bias and/or with suitable circuit geometry, up to the point that the coupling energy g can become comparable to the bare frequency of the NR element [303], leading to high sensitivity of the electromechanical set-up to phonon number.

In Fig. 3.8a we report the nonlinear shift $\delta = \omega_{21} - \omega_{10}$ obtained as a function of g via a numerical diagonalization of the Hamiltonian in Eq. (3.59). As typical proof-of-concept parameters we assumed $\omega_{NR} = 100$ MHz, $\Omega_{SC} = 500$ MHz, and a suitable number of bosonic excitations to obtain numerical convergence for the eigenvalues and eigenstates of interest. The target $\delta \simeq 1$ MHz is obtained for $g \simeq 50$ MHz, which is rather large but still accessible without too much effort [285–287, 303]. To compare with numerical results, we also report here the analytical expression of the




Figure 3.8: Non-linearity induced by superconducting circuit. (a) Nonlinear shift between the ground-first and the first-second excited states transitions $\delta = \omega_{21} - \omega_{10}$, induced by coupling a nanoresonator NR to a superconduction non-linear circuit SC. Results are obtained diagonalizing Eq. (3.59). (b) Probabilities of Fock states (with SC in $|0\rangle$, left axis) and of finding the SC in the excited state $|1\rangle$ (right axis) for the first three eigenvectors $|\psi_i\rangle$ of the Hamiltonian in Eq. (3.59). (c) Dissipation time scale T_1 for the coupled NR-SC system, obtained from Bloch-Redfield simulations. (d) Decoherence time scale T_2 for the coupled NR-SC system. The values of overall T_1 and T_2 timescales in (c)-(d) are expressed in units of the corresponding typical timescales for the superconducting circuit assumed in the simulation ($T_{1,SC} = 1$ ms and $T_{2,SC} \simeq 10 \mu\text{s}$). In all panels $\omega_{NR} \equiv \omega = 100$ MHz and $\Omega_{SC} \equiv \Omega = 500$ MHz

nonlinear shift as obtained to fourth order perturbation theory in the SC-NR oscillators coupling

$$\delta \simeq 2g^4 \left[\frac{2}{(\Omega - \omega)^2(\Omega + \omega)} + \frac{2}{(\Omega + \omega)^2(\Omega - \omega)} + \frac{1}{(\Omega + \omega)^3} + \frac{1}{(\Omega - \omega)^3} \right] \quad (3.60)$$

Here, we used the compact notation $\Omega_{SC} \equiv \Omega$ and $\omega_{NR} \equiv \omega$. Moreover, all terms but the last one are only present if the effect of counter-rotating terms in Eq. (3.59) is taken into account.

To better unveil the actual character of the collective excitations in the first three energy levels, ψ_i (i.e. the designated computational basis plus one extra level), we report in Fig. Fig. 3.8b the amplitude probability of the bare Fock eigenstates of the uncoupled harmonic oscillator on the corresponding new eigenstate of the coupled system, $p_n = |\langle n,0|\psi_n\rangle|^2$. Here the second index in the bra vector indicates the SC in its ground state, $|0\rangle$. On the right axis of the same Fig. 3.8b we show the total probability, for each collective eigenstate, to find the SC in the $|1\rangle$ excited state, regardless of the state of the NR: this information gives an estimate of the amount of wavefunction leaking on the SC as a consequence of the coupling, that is the magnitude of the mixing of the bare degrees of freedom. As it can be seen, in the region of interest such mixing never exceeds $\sim 5\%$

for the relevant states in the computational basis. On one hand, this last result prevents us from considering this combined NR-SC architecture as a form of full hybrid encoding: indeed, given the dispersive nature of the coupling, the two subsystems are not treated on equal footing, with the NR degrees of freedom maintaining a predominant role. On the other hand, such regime is crucial in guaranteeing that the NR performances in terms of coherence are not significantly affected by the presence of the additional SC element.

We quantitatively analyzed the actual effect of introducing the SC element in terms of dissipation (T_1) and coherence times (T_2) of the collective excitations by applying the formalism of open quantum systems. In particular, given the strong internal coupling (as we saw above, g can get up to a significant fraction of ω_{NR}) and of the presence of non-negligible counter-rotating terms, we employed a Bloch-Redfield master equation with suitable secular approximation [296]. We start by diagonalizing the full system Hamiltonian, a procedure that is ultimately consistent with the fact that the actual qubits are defined as slightly mixed excitations, and then derive the dissipation and pure dephasing terms by computing the matrix elements of a set of operators describing the coupling to the environment. We considered an extended system

$$H_{TOT} = H_{NR+SC} + H_{env,NR} + H_{env,SC} + H_I \quad (3.61)$$

where H_{NR+SC} is given in Eq. (3.59), while

$$H_{env,i} = \sum_k \omega_{k,i} d_{k,i}^\dagger d_{k,i} \quad (3.62)$$

is a collection of harmonic bath modes ($i = NR, SC$) and the system-bath coupling H_I is a sum of terms of the form

$$H_{I,S} = O \otimes \sum_k g'_{k,i} (d_{k,i} + d_{k,i}^\dagger). \quad (3.63)$$

Here O is a hermitian system operator describing individual interaction mechanisms for either the nanoresonator or the superconducting circuit. We use, for example, the operators $O_{NR} = b + b^\dagger$ and $O_{SC} = \sigma_- + \sigma_+$ to describe what in the bare SC and NR basis would be interpreted as dissipative processes, while coupling via $O_{NR} = b^\dagger b$ and $O_{SC} = \sigma_z$ gives rise to dephasing contributions. The resulting master equation will fully describe all possible transitions induced by the environment on the effective dynamics of the coupled nanoresonator and SC system, with rates that are proportional to the spectral functions of the environment evaluated at the relevant transition frequencies. Such functions are taken for simplicity as zero temperature white noise spectra

$$S_i(\omega) = \begin{cases} \gamma_{i,d} & \text{for } \omega = 0 \\ \gamma_i & \text{for } \omega > 0 \\ 0 & \text{for } \omega < 0 \end{cases} \quad (3.64)$$

where, again, $i = NR, SC$ and γ (γ_d) represent dissipation (pure dephasing) contributions in the uncoupled case. These numerical simulations were carried out with the QuTiP library in Python [204, 205]. In Fig. 3.8c-d we show the change in the total T_1 (decay of diagonal terms) and T_2 (decay of coherences) of the coupled system as a function of g , obtained by observing the time evolution of an initial superposition of the computational basis elements. More explicitly, data are obtained by fitting the exponential decay of the excited state occupation probability and of the off-diagonal element (coherence) of the resulting density matrix. The following parameters were used in these simulations: $\omega_{NR} = 100$ MHz, $\Omega_{SC} = 500$ MHz, $\gamma_{NR} = 50$ Hz (corresponding to $T_{1,NR} = 20$ ms), $\gamma_{NR,d} = 200$ Hz (i.e. $T_{2,NR} = 8$ ms), $\gamma_{SC} = 1$ kHz, and $\gamma_{SC,d} = 50$ kHz. In particular, dissipation rates for the low-frequency SC element correspond to $T_{1,SC} = 1$ ms and $T_{2,SC} \simeq 10 \mu\text{s}$, as experimentally reported in Ref. [302]. Our results clearly predict that the additional superconducting element affects the original performances of the nanomechanical oscillator by less than an order of magnitude, thus still preserving a significant advantage over the typical dissipation and coherence times of the SC element alone, whose T_1 and T_2 values are taken as a reference for normalization in the plot.

3.3.5 Digital quantum simulations

The promising theoretical predictions for the fidelities of elementary single- and two-qubit gates suggest that the proposed electromechanical architecture could potentially meet the requirements for building a digital quantum simulator, capable of running long sequences of concatenated gates. In accordance to the framework introduced in Chapter 1, we will now briefly describe the realization of a few prototypical algorithms for the digital quantum simulation of spin models, using the fundamental set of operations accessible to the our electromechanical quantum computing set-up. At difference with Chapters 1 and 2, we focus here on the hardware performances, more than on the algorithm design: indeed, after decomposing the target Hamiltonian in the corresponding sequence of gates using standard techniques, we numerically simulated the actual experimental procedure to operate the electromechanical qubits, including all required external pulses and realistic dissipation and decoherence processes. This in turn means integrating Eq. (3.28), with $H_{\text{control}}(t)$ being the translation of a given digital quantum simulation algorithm into the corresponding series of time-dependent frequency shifts and driving pulses needed to activate the quantum operations, as described in Sec. 3.3.3. All the results were obtained assuming $\gamma_{TR,d}/2\pi = 100$ kHz as realistic transmon dephasing rate and for different choices of the nanomechanical qubits T_2 times.

As a first example, we consider the problem of spin $S = 1$ total magnetization tunneling, previously discussed in Sec. 2.2.1. We recall here for convenience that the physical model is described by the Hamiltonian

$$\mathcal{H} = DS_z^2 + E(S_x^2 - S_y^2) \quad (3.65)$$

which can be decomposed, by writing S_α as the sum of two spin-1/2 operators, as

$$H = 2Ds_z^{(1)}s_z^{(2)} + 2E\left(s_x^{(1)}s_x^{(2)} - s_y^{(1)}s_y^{(2)}\right) \quad (3.66)$$

Following Sec. 2.2.1, the time evolution of the total magnetization $S_z(t) = s_z^{(1)}(t) + s_z^{(2)}(t)$ can be simulated on a quantum register of two qubits by realizing a sequence of gates representing the unitary operator $U(t) = e^{-iH_E t}$ generated by

$$H_E = \frac{E}{2}\left(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)}\right) \quad (3.67)$$

Using the native $U_{XY}(t)$ two qubit operation, see Eq. (3.50), it is straightforward to check that the quantum circuit in Fig. 3.9a achieves the desired result. Indeed, we can verify by direct inspection that

$$e^{-i\frac{E}{2}\left(\sigma_x^{(1)}\sigma_x^{(2)} - \sigma_y^{(1)}\sigma_y^{(2)}\right)t} = R_x^{(1)}(\pi)e^{-i\frac{\Gamma}{8}\left(\sigma_x^{(1)}\sigma_x^{(2)} + \sigma_y^{(1)}\sigma_y^{(2)}\right)t}R_x^{(1)}(-\pi) \quad (3.68)$$

if we assume⁸ $E = \Gamma/4$. At the hardware level, the algorithm requires two gaussian pulses on qubit 1 to implement $R_x^{(1)}(\pi)$ rotations and frequency shifts on the two qubits and the intermediate transmon to activate the $U_{XY}(t)$ exchange evolution, the duration of the latter varying with the target evolution phase $\lambda \propto Et$: a schematic representation is provided in Fig. 3.9a below the quantum circuit. The resulting time evolution of the total magnetization obeys, given an initial state $|\uparrow\uparrow\rangle$,

$$S_z(t) = s_z^{(1)}(t) + s_z^{(2)}(t) = \cos(2Et) \quad (3.69)$$

and no Trotter steps are needed in this case. Numerical results are reported in Fig. 3.9a, where we compare the exact result of Eq. (3.69) to the digital quantum simulation performed on the electromechanical set-up, for either $\gamma_{NR,d} = 0$ or $\gamma_{NR,d}/2\pi = 1$ kHz. The overall fidelity, averaged over all points, equals $\mathcal{F} = 0.999$ for $\gamma_{NR,d} = 0$ and $\mathcal{F} = 0.988$ for $\gamma_{NR,d}/2\pi = 1$ kHz, respectively.

A slightly more challenging digital simulation is reported in Fig. 3.9b, where the target model is now the two-spin Heisenberg Hamiltonian, see also Eq. (1.83)

$$H_{\text{Heis},2} = J\left(\sigma_x^{(1)}\sigma_x^{(2)} + \sigma_y^{(1)}\sigma_y^{(2)} + \sigma_z^{(1)}\sigma_z^{(2)}\right) \quad (3.70)$$

⁸Typically, Γ is a well defined property of the hardware: arbitrary values of the target E parameter can then be obtained for example by suitably rescaling the physical phase acquired during a gate operation with respect to the simulated time in $U(t) = e^{-iH_E t}$.




Figure 3.9: Digital quantum simulations with electromechanical qubits. (a) Tunneling of $S = 1$ total magnetization. A scheme of the pulse sequence operated at hardware level is shown below the quantum circuit (drive amplitude and frequency shifts not to scale). (b) Heisenberg model on two spin-1/2, with solid lines indicating exact results and markers indicating results from the numerical simulation of the electromechanical setup. No Trotter approximation is required. (c) Transverse Field Ising model for $n = 10$ Trotter steps and different choices of the coherence time for the nanoresonators. In all panels, we use the compact notation $U_{XY}(\delta) = \exp[-i\delta(\sigma_x^{(1)}\sigma_x^{(2)} + \sigma_y^{(1)}\sigma_y^{(2)})]$.

The quantum algorithm is now designed, following Fig. 1.4c and the identity [53, 128]

$$e^{-iH_{\text{Heis},2t}} = e^{-i\frac{J}{2}H_{xxyy}t} e^{-iH_{xxxx}t} e^{-iH_{zzyy}t} \quad (3.71)$$

where $H_{\alpha\alpha\beta\beta} = \sigma_\alpha^{(1)}\sigma_\alpha^{(2)} + \sigma_\beta^{(1)}\sigma_\beta^{(2)}$. For $\gamma_{NR,d}/2\pi = 0.1$ kHz (i.e. $T_{2,NR} = 10$ ms) and $J = \Gamma/8$, the average fidelity of the simulation is above 99%.

Finally, the digital quantum simulation of the 2-spin Transverse Field Ising model (see Eq. (1.91))

and Sec. 2.2.3)

$$H_{\text{TIM},2} = \Lambda s_x^{(1)} s_x^{(2)} + b(s_z^{(1)} + s_z^{(2)}) = \frac{\Lambda}{4} \sigma_x^{(1)} \sigma_x^{(2)} + \frac{b}{2} (\sigma_z^{(1)} + \sigma_z^{(2)}) \quad (3.72)$$

gives the opportunity to test longer sequences of gates. Indeed, in this case a non-trivial application of the Trotter approximation formula is required, leading to the following construction

$$e^{-iH_{\text{TIM},2}t} \simeq \left(R_z^{(1)}(b\tau) R_z^{(2)}(b\tau) U_{\text{XY}} \left(\frac{\Lambda}{\Gamma} \tau \right) R_x^{(1)}(-\pi) U_{\text{XY}} \left(\frac{\Lambda}{\Gamma} \tau \right) R_x^{(1)}(\pi) \right)^n \quad (3.73)$$

where $\tau = t/n$ and, using Eq. (3.50) for the definition of the native U_{XY} in terms of hardware parameters, we made explicit the possible rescaling⁹ of the physical phase to be acquired by the electromechanical qubits when in general $\Lambda \neq \Gamma$. The quantum circuit applied on the hardware is shown in Fig. 3.9c: measuring the state of the electromechanical qubits gives direct access to the total magnetization along the x axis¹⁰, i.e. $\langle S_x \rangle = \text{Tr}[\rho(s_x^{(1)} + s_x^{(2)})]$. Numerical simulations reported in Fig. 3.9c were done choosing $\Lambda = 2b = \Gamma$ and $n = 10$ Trotter steps. For each data point, this requires the sequential concatenation of 20 two-qubit and 40 single-qubit gates (the latter can often be operated in parallel on both qubits). The digital simulation is therefore quite demanding in terms of hardware performances: nevertheless, the fidelity is already 0.90 for $\gamma_{NR,d}/2\pi = 1$ kHz, and it increases to 0.96 by assuming a more optimistic value $\gamma_{NR,d}/2\pi = 100$ Hz. The results are especially noteworthy if we consider that the total physical time required to operate the longest sequence of gates, corresponding to the last point in Fig. 3.9c, is about $150 \mu\text{s}$. The overall potential of the proposed electromechanical architecture for quantum information processing and quantum simulation purposes is thus confirmed.

To summarize, in this chapter we have briefly explored a few possible future directions for quantum computing technology. While trapped ions and superconducting circuits continue their steady approach to maturity and receive attention from a growing academic and industrial community, the many open challenges in quality, control and scalability make the search for new paradigms as exciting and scientifically relevant as the improvement of existing solutions. Here we have proposed quantum electromechanics as a promising candidate platform, theoretically describing a novel architecture for quantum information processing based on state-of-the-art nano-devices. We demonstrated how qubits can be encoded in the anharmonic vibrational modes of mechanical resonators and integrated with superconducting circuitry. In our proposed set-up, switchable coupling mediated by virtual transmon excitations can be used, in combination with single-qubit manipulation techniques, to engineer a universal set of elementary quantum gates, whose fidelity remains remarkably high even with the inclusion of realistic decoherence and dissipation rates on both the nanoresonators and the superconducting elements. Finally, we concatenated such elementary operations to numerically assess the potential performances in the realization of paradigmatic quantum simulation algorithms, finding very encouraging results.

⁹Notice that, even though this is not explicit in Eq. (3.73), this rescaling can be most effectively done taking total phases modulo 2π , in order to keep the physical simulation time under control.

¹⁰Notice that, with respect to Chapter 1 and 2, here we used an equivalent parametrization of the TIM where the external field is along the z direction.

Quantum artificial neural networks

As quantum technologies start leaving academic research laboratories and set the stage for future commercialization, great effort is being devoted to harness the potential advantage of quantum information processing and employ it outside its traditional field of application. Even though the first practical demonstrations of quantum speedup over classical counterparts will probably come from the solution of rather abstract problems [81] or within very special settings [304, 305], a plethora of practical use cases for quantum computing is being explored and developed, from communication to optimization, from chemistry [180] to finance [87, 88]. The aim of such attempts, for which at present only proof-of-principle experimental tests on small noisy quantum processors can be given, is, on one hand, to develop the algorithmic and software machinery enabling future full-scale realizations, while, on the other hand, fulfilling a roadmap designed to bring the performances of near term quantum technology beyond classical counterparts in real world scenarios. Both sides require fundamental research in quantum computer science and a clear interdisciplinary attitude.

Among the fields that have been recently associated to quantum computing, machine learning attracted special attention. Indeed, it is certainly tempting and fascinating to put together two approaches that, in parallel and within the respective areas, have been holding promise of a modern revolution in computation and information science. We will devote the first part of this chapter to a brief account of the birth and status of Quantum Machine Learning (QML) as a discipline, laying out proposed advantages and current bottlenecks. We will then get back to the roots of machine learning, namely to Rosenblatt's and McCulloch-Pitts' artificial neuron models, and describe a series of original works [306, 307] in which such fundamental ingredients and their combinations have been translated into quantum procedures and employed to carry out simple classification and recognition tasks on real prototype quantum processors.

4.1 The rise of quantum machine learning

Following the growing interest and enthusiasm towards the so called artificial intelligence in many areas of human knowledge, machine learning techniques have successfully been applied to physical problems [308], ranging from cosmology and sub-atomic particles to many-body and quantum physics. In these approaches, classical computing architectures are employed for example to analyze and classify quantum states, to optimize experimental set-ups or quantum circuit design. In its purest form, Quantum Machine Learning (QML) [89, 309–311] aims at turning the relationship around, namely at developing quantum computing algorithms and procedures to analyze complex series of either classical or, in perspective, quantum data.

4.1.1 Quantum subroutines for linear algebra

The most natural point of contact between quantum information processing and machine learning is probably linear algebra. Indeed, many of the most demanding protocols in artificial intelligence and data analysis feature matrix manipulations in high-dimensional vector spaces, while, at the same

time, many renowned quantum algorithms precisely tackle some fundamental algebraic tasks, such as Fourier transformation [5, 312], matrix diagonalization [313] and the solution of linear systems of equations [314]. The whole corpus of quantum algorithms for linear algebra applications is already sufficiently mature that it is often identified with the acronym q-BLAS [89], for quantum Basic Linear Algebra Subroutines. In many cases, there exist theoretical proofs showing an exponential advantage of q-BLAS procedures over classical counterparts, many of them based on the intimate linear structure of quantum mechanics associating a 2^N -dimensional complex vector space to the joint state of N qubits. As an example, let us consider the forefather of many recent advances in the field, namely the Harrow-Hassidim-Lloyd (HHL) algorithm [314] for linear systems of equations. The problem solved by the HHL subroutine is indeed a classic in linear algebra: given a $D \times D$ Hermitian matrix A and a vector \vec{b} , find a vector \vec{x} such that

$$A\vec{x} = \vec{b} \quad (4.1)$$

The algorithm can be sketched as follows. First, the vector \vec{b} is, without loss of generality, normalized and written onto the quantum register of approximately $\log_2 D$ qubits as a superposition of computational basis states, e.g. $\vec{b} \mapsto |b\rangle = \sum_i b_i |i\rangle$. Then, by applying the well known phase estimation algorithm [5, 126], essentially involving the digital quantum simulation of the unitary operator e^{-iAt} , one can construct the state

$$|\psi\rangle = \sum_j \beta_j |u_j\rangle |\lambda_j\rangle \quad (4.2)$$

where u_j are the eigenvectors of A , the β_j are the coefficients of the (always possible, in principle) expansion of $|b\rangle$ on the basis $\{|u_j\rangle\}$ and $|\lambda_j\rangle$ are the eigenvalues of A written on an auxiliary quantum register. Adding an additional qubit and performing a rotation conditioned on $|\lambda_j\rangle$ we obtain

$$|\psi\rangle|0\rangle \mapsto \sum_j \beta_j |u_j\rangle |\lambda_j\rangle \left(\sqrt{1 - \frac{C}{\lambda_j}} |0\rangle + \frac{C}{\lambda_j} |1\rangle \right) \quad (4.3)$$

where C is a normalization constant. After uncomputing the phase estimation, we can discard the auxiliary quantum register on which the eigenvalues of A were written, leaving the state

$$\sum_j \beta_j |u_j\rangle \left(\sqrt{1 - \frac{C}{\lambda_j}} |0\rangle + \frac{C}{\lambda_j} |1\rangle \right) \quad (4.4)$$

Finally, a measure of the last qubit in the computational basis, conditioned on the outcome $|1\rangle$, puts the original quantum register in the state

$$\sqrt{\frac{1}{\sum_j C^2 |\beta_j|^2 / |\lambda_j|^2}} \sum_j \frac{\beta_j}{\lambda_j} |u_j\rangle \propto |A^{-1}b\rangle \equiv |x\rangle \quad (4.5)$$

The desired vector \vec{x} is thus recovered probabilistically, and the number of times the algorithm must be repeated to succeed depends on the condition number of the matrix A . For well behaved cases, the HHL subroutine runs in $O((\log D)^2)$ quantum operations, with an exponential advantage over classical counterparts requiring at least $O(D \log D)$ steps. A number of caveats must however be taken into account, mostly regarding loading of classical data (i.e. the vector b) on a quantum register and the efficient extraction of the output: indeed, a generic quantum state $|b\rangle$ can be exponentially expensive to prepare from scratch, and fully characterizing a state $|x\rangle$ to recover its components requires an exponential number of tomographic measurements. As we will see, the interface layer between classical data and their quantum representation is indeed a very common bottleneck in many quantum machine learning schemes: proposed solutions go from directly inputting quantum states stored in a quantum Random Access Memory (qRAM) [315] to post-processing output data directly on the quantum register to recover only partial information of interest, possibly encoded

into the expectation value of some quantum observable $m_x = \langle x|M|x\rangle$, though this is, in general, still an open line of research.

Another example, using quantum techniques similar to the HHL algorithm and with similar scaling properties, is the so called quantum Principal Component Analysis (qPCA) [316], which, for a series of data written in D -dimensional vectors \vec{v}_j , finds the eigenvectors associated to the largest eigenvalues (i.e. the principal components) of the covariance matrix $C = \sum_j \vec{v}_j \vec{v}_j^T$ exponentially faster than the best corresponding classical algorithm.

4.1.2 Quantum classification

Classification of data is one of the typical tasks supervised and unsupervised machine learning systems are programmed for. In particular, the well known problem of pattern classification can be expressed in the following terms [309]: given a training set $\mathcal{T} = \{(\vec{v}_i, c_i)\}$ containing feature vectors \vec{v}_i and their corresponding class c_i , with $c_i \in \{C_k\}_1^L$ if there are L possible classes, the algorithm must learn how to assign a new and unseen vector x to one of the classes, according to general rules inferred from \mathcal{T} . Several quantum classifiers have been recently proposed and tested [317–324], relying on different forms of encoding and data processing and sometimes mimicking different classical machine learning algorithms. Oracle-based quantum classification has also been put forward as a possible route to the demonstration of quantum advantage on noisy quantum processors [325, 326].

Estimating distances between feature vectors, thus translating in mathematical terms the common wisdom that similar properties are a proxy for group identity, is at the basis of some simple yet effective classification algorithms. For example, the k -nearest neighbors method constructs clusters of data by assigning a vector \vec{x} to the class c_x to which its k (already classified or assigned in \mathcal{T}) nearest vectors belong. A translation of this and similar approaches in the quantum computing paradigm focuses naturally on efficient evaluations of vector distances, which can be again reduced to an algebraic problem. For example, it has been proposed that the overlap (i.e. the fidelity) $|\langle a|b\rangle|$ between two quantum states $|a\rangle$ and $|b\rangle$ encoding for feature vectors can be used as a similarity measure. The so called SWAP test [327] is a way of computing such a quantity on a digital quantum computing architecture, and is based on the following quantum circuit:

$$(4.6)$$

where we used the notation

$$\text{SWAP} = \begin{array}{c} \times \\ \updownarrow \\ \times \end{array} \quad (4.7)$$

Indeed, it is straightforward to verify that the output state is

$$|\psi_{\text{ST}}\rangle = \frac{1}{2} [|0\rangle(|a\rangle|b\rangle + |b\rangle|a\rangle) + |1\rangle(|a\rangle|b\rangle - |b\rangle|a\rangle)] \quad (4.8)$$

and that the probability of observing the first qubit in state $|0\rangle$ is therefore directly linked to the desired overlap

$$p_0 = \frac{1}{2} + \frac{1}{2} |\langle a|b\rangle|^2 \quad (4.9)$$

Through the simple encoding onto computational basis states

$$\vec{x} \mapsto |x|^{-1/2} \sum_i x_i |i\rangle \quad (4.10)$$

the SWAP test has been used by Lloyd *et al.* [317] to construct a quantum version of the nearest centroid algorithm, namely a variation of the k -nearest neighbors in which distances are computed

between the vector to be classified and the “center of mass” of the classes in the feature space. With a slightly more involved form of encoding, Wiebe *et al.* [328] reported a quadratic speedup with respect to classical algorithms in computing the inner product of two vectors.

More recently, quantum-enhanced support vector machines (SVM) have been designed [321, 322] and realized on real superconducting quantum processors [321], showing promising performances. Support vector machines are used for linear discrimination, a method that searches for the best separating hyperplane discriminating between two class regions and acting as a boundary for classification of unknown inputs. The required linear separability of feature vectors belonging to different classes enables, on one hand, the use of standard mathematical techniques such as the method of Lagrange multipliers to simplify the solution of the hyperplane optimization problem. On the other hand, it would severely limit the applicability of the method to those cases in which the classes are geometrically disjoint. This restriction can be overcome by mapping a given problem into a higher dimensional feature space, namely by changing the encoding and the effective distance measure applied to the data: the fundamental quantity for SVMs applications therefore becomes the so called kernel K , containing information about all the inner products of mapped feature vectors. Indeed, given two such D -dimensional vectors \vec{x} and \vec{x}' , a valid kernel $K(\vec{x}, \vec{x}')$ induces a distance measure which serves as the basis to compute the separating hyperplane. Different kernels provide different metrics and, as a result, different outcomes of the classification problem. In quantum SVMs [322], one usually defines a feature map $\phi : \mathcal{X} \mapsto \mathcal{F}$ relating classical data vectors $\vec{x} \in \mathcal{X}$ to quantum states $|\phi(\vec{x})\rangle$ living in the so called quantum feature Hilbert space \mathcal{F} . The corresponding kernel is thus constructed using the natural inner product in \mathcal{F} :

$$K(\vec{x}, \vec{x}') = \langle \phi(\vec{x}) | \phi(\vec{x}') \rangle \quad (4.11)$$

Notice that, in general, the map ϕ can contain non-linear combinations of the input components. Quantum information processing techniques can then be used to build kernels which may be hard to compute classically or to speed up the evaluation of inner products between feature vectors. Rebentrost *et al.* [329] have for example proposed that, through the direct encoding defined in Eq. 4.10, a standard kernel $K(\vec{x}, \vec{x}') \propto \vec{x} \cdot \vec{x}'$ can be computed efficiently via a partial trace operation, i.e. by discarding part of the information stored in the composite quantum register used for training and/or loading the data. Indeed, if $\{|x_i\rangle\}$ is a basis of the 2^N -dimensional training vector space \mathcal{T} , in such a way that any quantum state representing training vectors \vec{v}_j can be expressed as $|v_j\rangle = \sum_i \alpha_{ji} |x_i\rangle$, one can first use two N -qubit quantum registers to prepare the state

$$|\chi\rangle = \frac{1}{\sqrt{N_X}} \sum_{i=1}^{2^N} |\vec{x}_i||i\rangle|x_i\rangle \quad (4.12)$$

where $N_X = \sum_{i=1}^{2^N} |\vec{x}_i|^2$ and $|i\rangle$ are computational basis states. Then, using Eq. 4.10 it is easy to see that

$$\langle x_i | x_j \rangle = \frac{\vec{x}_i \cdot \vec{x}_j}{|\vec{x}_i||\vec{x}_j|} \quad (4.13)$$

and therefore

$$\text{Tr}_X [|\chi\rangle\langle\chi|] = \frac{1}{N_X} \sum_{i,j=1}^{2^N} |\vec{x}_i||\vec{x}_j| \langle x_j | x_i \rangle |j\rangle\langle i| = \frac{\hat{K}}{\text{Tr}[\hat{K}]} \quad (4.14)$$

where Tr_X denotes the partial trace over the quantum register storing the $|x_i\rangle$ states and $(\hat{K})_{ij} = K(\vec{x}_i, \vec{x}_j)$. In the work by Havlíček *et al.* [321], quantum feature maps were put into practice via digital quantum circuits. As an example, the authors report the following sequence of parametrized quantum operations preparing a class of quantum states $|\phi(\vec{x})\rangle$ of N qubits whose inner product is conjectured to be hard to compute classically:

$$\mathcal{U}_{\phi(\vec{x})} = \exp\left(i \sum_S \phi_S(\vec{x}) \prod_{i \in S} \sigma_z^{(i)}\right) \text{H}^{\otimes N} \exp\left(i \sum_S \phi_S(\vec{x}) \prod_{i \in S} \sigma_z^{(i)}\right) \text{H}^{\otimes N} \quad (4.15)$$

Here S are subsets of qubit indexes $1, \dots, N$ and a possible choice of $\phi_S(\vec{x})$ for $|S| \leq 2$ is

$$\phi_i(\vec{x}) = x_i \quad \phi_{\{i,j\}}(\vec{x}) = (\pi - x_i)(\pi - x_j) \quad \text{for } i, j = 1, 2 \quad (4.16)$$

Following Eq. 1.48, it is not difficult to see that the elementary building block of the algorithm can be decomposed in terms of elementary entangling operations and single qubit rotations

$$e^{-i\phi_{\{i,j\}}(\vec{x})\sigma_z^{(i)} \otimes \sigma_z^{(j)}} = \begin{array}{c} \bullet \text{---} \text{---} \text{---} \bullet \\ | \quad \quad | \\ \oplus \text{---} \boxed{R_z(2\phi_{\{i,j\}}(\vec{x}))} \text{---} \oplus \end{array} \quad (4.17)$$

while single qubit R_z gates are sufficient when S contains only one index. With these techniques, a 2-qubit variational classification protocol was then tested on real IBM Q quantum processors, obtaining very high accuracy on simple proof of concept experiments. Finally, Schuld and Killoran [322] proposed squeezing in bosonic Fock space as a way of implementing a feature map in photonic continuous variable quantum architecture.

4.1.3 Searching for quantum neurons

In the early days of machine learning, artificial neurons and neural networks were first proposed as powerful tools for trainable classification and pattern recognition [330, 331]. Gaining momentum after a few milestones results such as the backpropagation algorithm for training [332] and the Universal Approximation Theorem [333, 334], deep feed-forward neural networks experienced rapid evolution and development, and constitute today the cornerstone of many artificial intelligence protocols [335, 336].

The simplest “digital” formalization of an artificial neuron was provided by McCulloch and Pitts [330]. Every such node is represented by a binary variable $X_i \in \{-1, 1\}$, with $X_i = 1(-1)$ denoting an active (rest) state. Every neuron is controlled by a set of inputs x_1, \dots, x_m which can either be signals from other neurons in the network (i.e. $x_j \equiv X_j$) or external data. In every computational step, each input is first weighted by a synapse coefficient $w_{ij} \in \{-1, 1\}$. Then, the state of the i -th neuron is updated via an integrate-and-fire response

$$X_i^{\text{updated}} = \begin{cases} 1 & \text{if } \sum_j w_{ij}x_j \geq \theta_i \\ -1 & \text{otherwise} \end{cases} \quad (4.18)$$

where θ_i represents some predefined threshold. A full neural network constructed by integrating several McCulloch-Pitts nodes can thus store binary patterns $\{X_1, \dots, X_N\}$ in the collective activation state. The dynamics of the network can either proceed in sequential order (as it is usually the case in feed-forward neural networks), synchronous or random order. A particularly interesting case is provided by Hopfield neural networks (HNNs) [337], which display remarkable capabilities of reproducing the behavior of associative memory. The latter is the ability of recovering a previously stored item by relying on some partial information about it: for example, the letters *Einst* immediately point to Albert Einstein in an archive of notable physicists. HNNs (see Fig. 4.1a) are in principle symmetric and fully connected, i.e. $w_{ij} = \tilde{w}_{ij}(1 - \delta_{ij})$ with $w_{ij} = w_{ji}$, and are able to store a set of specific binary patterns and retrieve them afterwards. Indeed, let us denote by \vec{V}^s a particular state of the N sites, namely

$$\vec{X}^s = \{X_1^s, X_2^s, X_3^s, \dots, X_N^s\} \quad (4.19)$$

It can be shown that, by choosing synapses according to the so called Hebbian rule

$$w_{ij} = \sum_s (2X_i^s - 1)(2X_j^s - 1) \quad (4.20)$$

for some collection of possible states \vec{X}^s , then all these states are stable under the updating algorithm for neuron states. The existence of such stable limit points is further confirmed by the fact that the energy functional

$$E = -\frac{1}{2} \sum_{i \neq j} w_{ij} X_i X_j + \sum_i \theta_i X_i \quad (4.21)$$




Figure 4.1: Artificial neural networks. (a) Example of the typical layout for a Hopfield neural network. All nodes are equivalent and connected to all others. (b) Scheme of a feed-forward neural network, with one hidden layer surrounded by input and output ones.

is monotonically decreasing under random updates of the neurons [337]. The state of the system will then continue to change until a (local) minimum of E is reached: content-addressability finally derives from the fact that, by providing a certain initial state $\{X_1, \dots, X_N\}$, the system converges to the \vec{X}^s , among the ones encoded in the weights, which is closer to the input in terms of Hamming distance (i.e. the number of differing bits, or ± 1 entries). Of course, the number of different patterns which can reliably be stored and retrieved depends on the overall size and capacity of the network [338]. As Hopfield himself noted, there exist a clear isomorphism between HNNs and generalized Ising models where w_{ij} are interpreted as the exchange couplings.

A more refined artificial neuron model, proposed by Rosenblatt [331] and later further generalized, relaxes the constraint on the binary nature of state variables and is often referred to as “perceptron”. The output of such graded-response neurons now becomes continuous valued, for example $X_i \in [-1, 1]$, and is computed from weighted inputs using the output states of connected neurons (or other input signals) from a previous step:

$$X_i^{\text{updated}} = f \left(\sum_j w_{ij} X_j \right) \quad (4.22)$$

A common choice for the response function f is the sigmoid

$$f(x) = \frac{1}{1 + e^{-x}} \equiv \text{sgm}(x) \quad (4.23)$$

while binary-response neurons are recovered by choosing $f(x) = \text{sign}(x)$. When used in feed-forward neural networks, i.e. featuring a series of successive neuron layers with no feedback or recursive structures (see Fig. 4.1b), graded-response neurons can be trained adjusting the connections w_{ij} throughout the network with the well known backpropagation procedure [332], thus making feed-forward neural networks a very powerful and versatile tool to analyze, classify and manipulate data. As it is common in supervised learning, a training set is initially provided as an ensemble of pre-labeled input-output pairs $(\vec{x}, \vec{o}_{\text{target}})$. Here \vec{x} is a vector whose number of entries matches the number of inputs the network can accept and the dimension of \vec{o}_{target} is that of the typical output of the network (e.g. \vec{o}_{target} is N -dimensional if the network has N neurons in the output layer). Inputs are passed to the first layer of the neural network, and the states of all neurons are updated in a cascaded fashion using Eq. (4.22). After each input has gone through the network (typically initialized with random connections w_{ij}) the output \vec{o} is compared with the desired one $\vec{o}_{\text{target}} \equiv \vec{y}$: in particular, the following quadratic error function is constructed

$$\mathcal{E} = \frac{1}{2} \sum_k (y_k - o_k)^2 \quad (4.24)$$

A minimization of \mathcal{E} via gradient descent then gives the update rule for the weights. If, for a given neuron, we define $\chi_i = \sum_k w_{ik} X_k$ as the weighted sum of its inputs, we have

$$\frac{\partial \mathcal{E}}{\partial w_{ij}} = \frac{\partial \mathcal{E}}{\partial X_i} \frac{\partial X_i}{\partial \chi_i} \frac{\partial \chi_i}{\partial w_{ij}} \quad (4.25)$$

Let us now assume that we are working with a neuron (here labeled with i) in the output layer. Then $X_i = o_i$ and we get

$$\begin{aligned} \frac{\partial \chi_i}{\partial w_{ij}} &= X_j \\ \frac{\partial X_i}{\partial \chi_i} &= f'(\chi_i) \\ \frac{\partial \mathcal{E}}{\partial X_i} &= \frac{\partial \mathcal{E}}{\partial o_i} = -(y_i - o_i) \end{aligned} \quad (4.26)$$

As a result

$$\frac{\partial \mathcal{E}}{\partial w_{ij}} = -(y_i - o_i) f'(\chi_i) X_j \equiv -\delta_i X_j \quad (4.27)$$

and the weight w_{ij} should then be updated by a quantity $\Delta w_{ij} \propto \delta_i X_j$. In case the i -th neuron is part of a hidden layer, the evaluation of $\partial \mathcal{E} / \partial X_i$ is slightly more elaborated. Making again use of the chain rule for derivatives, we get

$$\frac{\partial \mathcal{E}}{\partial X_i} = \sum_{l \in L} \left(\frac{\partial \mathcal{E}}{\partial X_l} \frac{\partial X_l}{\partial \chi_i} \frac{\partial \chi_l}{\partial X_i} \right) \quad (4.28)$$

where L is the set of all neurons in the successive layer receiving an input from neuron i . Finally,

$$\frac{\partial \mathcal{E}}{\partial X_i} = \sum_{l \in L} \left(\frac{\partial \mathcal{E}}{\partial X_l} f'(\chi_l) w_{li} \right) \quad (4.29)$$

and the desired quantity can be recovered as soon as the derivatives $\partial \mathcal{E} / \partial X_l$ are known. Since we have seen above that this is true for the output layer, all derivatives can be computed recursively moving backwards from the last layer to the first. Once more, the weights are then updated with

$$\Delta w_{ij} \propto \sum_{l \in L} \left(\frac{\partial \mathcal{E}}{\partial X_l} f'(\chi_l) w_{li} \right) f'(\chi_i) X_j \equiv \delta_i X_j \quad (4.30)$$

The constant of proportionality, called the learning rate η , must carefully be selected to avoid a too strong sensitivity on single items within the training set, which could result in oscillations around the minimum of \mathcal{E} , but without sacrificing too much in the speed of convergence. To sum up, after each cycle the weights are updated with the rule

$$w_{ij} \mapsto w_{ij} + \Delta w_{ij} = w_{ij} + \eta \delta_i X_j \quad (4.31)$$

where the definition of δ_i depends on whether the i -th neuron is found in a hidden or output layer, see Eq. (4.27) and Eq. (4.30). Alternative schemes sometimes enhance the convergence and stability by using

$$\Delta w_{ij}(t+1) = \eta \delta_i X_j + \alpha \Delta w_{ij}(t) \quad (4.32)$$

i.e. by making the update of a weight w_{ij} at cycle $t+1$ dependent on the correction applied at the previous step.

Although mathematically simple, the training of large neural networks is often computationally expensive in terms of memory and time resources. This, together with the hope of recognizing and generating atypical data patterns [89], motivated over the last few years intense activity in search of a quantum version of artificial neurons and neural networks [339]. Rebentrost *et al.* [340] have for example designed a quantum version of Hopfield networks where neurons are encoded

in the amplitudes of quantum states, leading to an exponential advantage in storage resources. Using the HHL subroutine, the authors then envision a way of efficiently operating the network as an associative memory system. At the single node level, some proposals associate a qubit to each neuron [341] or binary state variable [320] and unitary operations to weights. In this setup, quantum phases can for example be used to compute and store the net input $\sum_j w_{ij}x_j$, which is then retrieved with a quantum Fourier transform [342]. Although not exploiting in full the exponential size of quantum Hilbert spaces, this scheme has the significant advantage that inputs can easily be presented in quantum superpositions, leading to quantum perceptron states with no classical analog. As a complementary solution, we will present in Sec. 4.2 a different encoding scheme for McCulloch-Pitts neurons which, similarly to Ref. [340], makes use of amplitudes in quantum superpositions to enhance memory capacity. In terms of basic neuron operations, interesting work has also been devoted to reconcile the inherent linear structure of quantum mechanics with the threshold behavior of activation functions [343, 344], for example by resorting to repeat-until-success techniques [345].

Regarding the learning phase, quantum search algorithms have been proposed as a way of speeding up the training process, e.g. by exploiting the well known Grover method [346–348]. In parallel, some attempts to achieve fully trainable quantum feed-forward neural networks [349] and to provide the quantum analog of classical backpropagation algorithms have also been reported [350]. It is, however, fair to say that an established theory of quantum learning and a clear investigation of the effective practical potential of quantum neural networks, both in terms of computational efficiency and required quantum resources, is still the subject of ongoing research. With this in mind, NISQ-accessible proposals and experiments, both in terms of single neuron dynamics and quantum-classical training schemes such as the one that we will discuss in section 4.2.3 and 4.2.3, could certainly provide new insights and offer the opportunity to assess with practical use cases the maturity of the field.

4.1.4 Perspectives on Quantum Machine Learning

We conclude this short and inevitably partial overview of the field of quantum machine learning by recalling a few other possible approaches beyond classification and supervised learning schemes.

Following the impressive results brought about by classical generative models [351], a number of promising strategies to employ similar methods in a quantum setting have been proposed. In particular, variational quantum autoencoders fit nicely into the current stream of heuristic and semi-heuristic approaches which seem particularly well suited for noisy intermediate-scale applications. Classical autoencoders can be trained to extract a low-dimensional representation (in the so called latent space) of higher-dimensional data and, in reverse, constitute the basis for generating new plausible samples not originally included in the training set: similarly, quantum autoencoders have been described primarily as a tool to compress quantum states into a smaller Hilbert space [352], thus enabling a more effective use of quantum resources. Quantum autoencoders have been put in relations with approximate quantum adders [353] and a few feasibility studies have already been performed e.g. on superconducting quantum processors [354].

In unsupervised learning, data are searched for patterns without providing explicit examples through a learning phase. A particularly interesting example is provided by reinforcement learning protocols, which are often based on the learning agent paradigm: here, the computer is given a set of rules or objectives and operates according to a reward/penalty mechanism. Through a combination of supervised and reinforcement learning, classical systems have come to master complex human games such as Go [355]. Potential quadratic speedups with respect to classical counterparts have been predicted [356] for quantum agent-based and reinforcement learning models, which have then been explored either from the perspective of superconducting architectures [357] or in trapped ions set-ups [358]. Superconducting quantum processors have also been employed in proof-of-principle demonstrations of hybrid unsupervised clustering protocols [359].

Despite our focus on the digital, circuit-based quantum computing paradigm, it is finally worth mentioning that several machine learning tasks could find breeding ground into other complementary approaches to quantum information processing. For example, proposed applications of quantum annealing procedures range from training of deep neural networks [360] to Boltzmann machines [361],

for which some preliminary work has been reported on D-Wave's systems [362]. In parallel, as already mentioned in Chapter 1, continuous variable quantum computing is also actively investigated as a possible way of implementing quantum neural networks on photonic architectures [246].

Quantum machine learning is a relatively young research field growing quickly. However, a number of challenges and open problems remain to be faced in the forthcoming years [89]. Apart from the obvious progress in hardware, which will be essential for all future quantum technologies, and from the search of new algorithms specifically designed for intermediate-scale and near-term architectures, most QML techniques developed so far still suffer from a series of bottlenecks and caveats that need to be overcome. Among these, we have already mentioned the often dominating exponential costs for loading classical data onto quantum registers and the equally hard task of extracting the processed information stored in quantum states. On top of it, a general assessment of the costs of QML algorithms, in terms e.g. of circuit depth or number of qubits, is needed as a preliminary step to provide clear and fair benchmarking with respect to classical procedures. The latter is also typically complicated by the fact that the actual scaling properties of many heuristic methods, either classical or quantum, is not yet fully understood. On a more specific sector of QML, as we saw in Sec. 4.1.3, a unifying theory of quantum learning would certainly be welcomed as a major accomplishment in the quest for fully quantum neural networks [309, 339]. Finally, we cannot ignore that a change of paradigm could be waiting around the corner: indeed, there is increasing awareness of the fact that the true power of QML may ultimately reside in the analysis of quantum data, which is by definition a difficult or inaccessible task for classical systems. As very recent examples, we can mention the results by Wright and McMahon [363], arguing that the capacity of classically parametrized quantum neural networks will ultimately be bounded by the corresponding classical value, the quantum reservoir processor by Gosh *et al.* [364], inspired by classical reservoir computing and capable of telling apart entangled states from separable ones, and the spiking neural-like architecture described by Kristensen *et al.* [365], which is based on elementary quantum spin components and can again be used to classify quantum states.

4.2 A quantum artificial neuron

As we saw in Sec. 4.1.3, artificial neural networks are designed as a set of nodes, or neurons, and the corresponding set of mutual connections, whose architecture is naturally inspired by neural nodes and synaptic connections in biological systems [335, 366]. The resulting class of computational models can naturally be applied to tasks such as pattern recognition, image classification, and decision making [366]. So far, artificial neural networks have been mostly run as classical algorithms on conventional computers, but considerable interest has also been devoted to physical neural networks, i.e. neural networks implemented on dedicated hardware [336, 367–369]. However, we have seen that there are valid reasons to believe that quantum processors, and with them the intrinsic capabilities of quantum mechanical systems to manipulate large vectors and matrices, could be employed to achieve exponential advantages in either memory storage or processing power for neural networks [339, 340, 342, 347, 349, 370].

Here we design a quantum procedure to closely mimic the functionality of a binary valued McCulloch and Pitts [330] artificial neuron. The m -dimensional classical input and weight vectors will be encoded on the quantum hardware by using N qubits, where $m = 2^N$, thus fully exploiting the exponential advantage of quantum information storage, as already pointed out [318, 340]. For loading and manipulation of data, we will make use of an original protocol to generate multipartite entangled states based on quantum information principles [371] optimizing the quantum computational resources to be employed. We will then experimentally show the effectiveness of our approach by practically implementing a 2-qubit version of the algorithm on the IBM Q quantum processors (see Sec. 2.1). In particular, we will use a single artificial neuron to sort out simple patterns, such as vertical or horizontal lines, in black and white collections of pixels, and we will provide a numerical demonstration of a possible hybrid training scheme. This building block will then be used in Sec. 4.3 to realize a first example of feed-forward neural network running on real quantum hardware.

4.2.1 Quantum algorithm for binary-valued artificial neurons

We start the discussion by outlining the general idea behind our proposed quantum procedure for artificial neurons. Let us consider two binary input and weight vectors

$$\vec{i} = \begin{pmatrix} i_0 \\ i_1 \\ \vdots \\ i_{m-1} \end{pmatrix} \quad \vec{w} = \begin{pmatrix} w_0 \\ w_1 \\ \vdots \\ w_{m-1} \end{pmatrix} \quad (4.33)$$

with $i_j, w_j \in \{-1, 1\}$ and $m = 2^N$. The primary target of the quantum algorithm is to perform the inner product $\vec{i} \cdot \vec{w}$ and then, based on the outcome, provide an activation or rest signal: in our design, this is achieved through a sequence of quantum operations aimed at loading the classical information contained in \vec{i} , followed by a \vec{w} -controlled manipulation and a final quantum measurement mimicking the thresholding behavior. Notice that we explicitly include all the practical steps needed to interface classical and quantum information, without e.g. assuming pre-loaded input quantum states: in such a way, the overall computational costs will be fairly assessed, taking into account the often relevant costs associated to those parts of QML protocols, see Sec. 4.1.1 and 4.1.4.

The Hilbert space associated to a quantum register containing N qubits has total dimension $m = 2^N$. A simple way of encoding binary-valued \vec{i} and \vec{w} can then be found by making use of the relative quantum phases (i.e. factors ± 1 in our case) in balanced superpositions: we can for example define the states

$$\begin{aligned} |\psi_i\rangle &= \frac{1}{\sqrt{m}} \sum_{j=0}^{m-1} i_j |j\rangle \\ |\psi_w\rangle &= \frac{1}{\sqrt{m}} \sum_{j=0}^{m-1} w_j |j\rangle \end{aligned} \quad (4.34)$$

where, as usual, we label computational basis states with integers $j \in \{0, \dots, m-1\}$ corresponding to the decimal representation of the respective binary string, and we recall that $i_j, w_j \in \{-1, 1\} \equiv \{e^{i\pi}, e^{i0}\}$. It is not difficult to see that, under the encoding scheme of Eq. (4.34), the inner product between inputs and weights is contained in the overlap

$$\langle \psi_w | \psi_i \rangle = \frac{\vec{w} \cdot \vec{i}}{m} \quad (4.35)$$

Assuming that the quantum register starts in the blank state $|0\rangle^{\otimes N}$, the first step of the algorithm prepares the state $|\psi_i\rangle$ by performing a unitary transformation U_i such that

$$U_i |0\rangle^{\otimes N} = |\psi_i\rangle \quad (4.36)$$

In principle, any $m \times m$ unitary matrix having \vec{i} in the first column can be used to this purpose, and we will give explicit examples in the following. In a more general scenario, nothing prevents replacing this step with e.g. a direct call to a quantum memory [315], although, as mentioned above and for the sake of completeness, in the present analysis we prefer to maintain an explicit description of the loading stage.

Let now U_w be a unitary operator obeying the constraint

$$U_w |\psi_w\rangle = |1\rangle^{\otimes N} = |m-1\rangle \quad (4.37)$$

As before, any $m \times m$ unitary matrix having \vec{w} in the last row satisfies this condition. If we apply U_w after U_i , the overall N -qubits quantum state becomes

$$U_w |\psi_i\rangle = \sum_{j=0}^{m-1} c_j |j\rangle \equiv |\phi_{i,w}\rangle \quad (4.38)$$




Figure 4.2: Artificial neuron models. (a) Classical artificial neuron: an input array \vec{i} is processed with a weight vector \vec{w} to produce a linear, binary valued output function. (b) Schematic representation of the proposed quantum algorithm for the implementation of an artificial neuron on a quantum processor: the first two unitary operations prepare the input quantum state, $|\psi_i\rangle$, and project it onto the $|\psi_w\rangle$ direction, writing the outcome onto the $|11\dots 1\rangle$ component. The final outcome is then extracted via an ancillary qubit, which is eventually measured to evaluate the activation state.

Using Eq. (4.37), we then have

$$\begin{aligned} \langle \psi_w | \psi_i \rangle &= \langle \psi_w | U_w^\dagger U_w | \psi_i \rangle = \\ &= \langle m-1 | \phi_{i,w} \rangle = c_{m-1} \end{aligned} \quad (4.39)$$

and the the desired result $\vec{i} \cdot \vec{w} \propto \langle \psi_w | \psi_i \rangle$ is contained, up to a normalization factor, in the coefficient c_{m-1} of the final state $|\phi_{i,w}\rangle$. To make this information accessible, an ancilla qubit (a), initially set in the state $|0\rangle$, is used as the target of a multi-controlled NOT gate, where the roles of controls is assigned to the N encoding qubits. This leads to [5]

$$|\phi_{i,w}\rangle |0\rangle_a \rightarrow \sum_{j=0}^{m-2} c_j |j\rangle |0\rangle_a + c_{m-1} |m-1\rangle |1\rangle_a \quad (4.40)$$

Finally, a non-linear threshold activation behavior is obtained by performing a quantum measurement: indeed, an observation of the state of the ancilla in the computational basis produces the output $|1\rangle_a$ (i.e. an active state of the neuron) with probability $|c_{m-1}|^2$. Despite its simplicity, we will show with specific examples that this choice is already sufficient to carry out some paradigmatic tasks. We notice, however, that once the inner product information is stored on the ancilla our procedure is in principle compatible with more refined approaches to the realization of threshold functions on quantum registers [343–345]. We also mention explicitly that, within our proposed encoding scheme, both parallel and anti-parallel (\vec{i}, \vec{w}) pairs produce an activation of the artificial neuron, while orthogonal vectors always result in the ancilla being measured in the state $|0\rangle_a$. This is a direct consequence of the probability of a quantum outcome being a quadratic function, i.e. $|c_{m-1}|^2$ in the present case, at difference with classical perceptrons that can only be employed as linear classifiers in their simplest realizations. In view of using our quantum model for pattern and/or image recognition and classification, this feature turns out to be particularly convenient, as it allows to interpret a given pattern and its negative on equivalent footing. More formally, this intrinsic symmetry reflects the invariance of the encoding $|\psi_i\rangle$ and $|\psi_w\rangle$ states under the addition of a global -1 factor.

A comparison of the classical artificial neuron model with the outline of our quantum procedure is provided in Fig. 4.2.




Figure 4.3: Elementary gate decompositions. (a) Example of a multi-controlled sign flip block $SF_{4,6}$ adding a -1 factor in front of the $|0110\rangle$ component of a 4-qubit computational basis, making use of an extra qubit in the $\sqrt{2}|-\rangle = |0\rangle - |1\rangle$ state. In each block, the NOT gates should be operated on those qubits for which the state is $|0\rangle$ in the target component. A standard construction to implement the multi-controlled C^N NOT gate using $N - 1$ additional working qubits (wq) is also shown. (b) Textbook decomposition of a Toffoli gate using CNOT and single-qubit operations.

Direct sign flip approach

After describing the general philosophy behind our proposed realization of a quantum circuit-based artificial neuron, we now move to discuss the actual implementation of the unitary operations U_i and U_w . We will first outline a “brute force” application of successive sign flip blocks, which we will then compare to an alternative and more effective approach based on the generation of the so called hypergraph states [371].

We define a sign flip block $SF_{N,j}$ as the unitary transformation acting on the computational basis of N qubits in the following way:

$$SF_{N,j}|j'\rangle = \begin{cases} |j'\rangle & \text{if } j \neq j' \\ -|j'\rangle & \text{if } j = j' \end{cases} \quad (4.41)$$

In general, $SF_{N,j}$ is equivalent to a multi-controlled quantum operation between N qubits: for example, for any $N, m = 2^N$, a controlled Z operation between N qubits ($C^N Z$) is a well known quantum gate [5] realizing $SF_{N,m-1}$, while a single qubit Z gate acts as $SF_{1,1}$. The whole family of sign-flip blocks for N qubits can be realized using $C^N Z$ gates in combination with single qubit NOT gates (i.e. single bit flip operations). Indeed, we can write

$$SF_{N,j} = O_j (C^N Z) O_j \quad (4.42)$$

where

$$O_j = \bigotimes_{l=0}^{N-1} (\text{NOT}_l)^{1-j_l} \quad (4.43)$$

In the expression above, NOT_l means that the bit flip is applied to the l -th qubit and $j_l = 0(1)$ if the l -th qubit is in state $|0\rangle(|1\rangle)$ in the computational basis state $|j\rangle$. Alternatively, the same result can also be obtained by using an extra ancillary qubit prepared in state $\sqrt{2}|-\rangle = |0\rangle - |1\rangle$ and multi-controlled NOT gates (C^NNOT), i.e. bit flip operations conditioned on the state of some control qubits. The latter decomposition is shown in Fig. 4.3a.

In view of practically testing the algorithm on real IBM Q processors, which feature only pairwise CNOT gates as native multi-qubit operations, we recall here that all N -controlled unitaries can be reduced to the corresponding single-control unitary operation with the addition of Toffoli gates


(4.44)

with matrix representation

$$\text{Toffoli} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \quad (4.45)$$

and by using $N - 1$ extra working qubits [5]. An example of this construction for a C^NNOT block is shown in Fig. Fig. 4.3a. The same can of course be applied to a C^NZ operation, reducing it to a $\text{C}^2\text{Z} \equiv \text{CZ}$


(4.46)

An IBM Q-accessible sequence can then be obtained via the identity


(4.47)

and the standard exact decomposition of the Toffoli gate in terms of CNOT and single qubit rotations [5, 91] reported for completeness in Fig. 4.3b using the shorthand notation

$$\text{T} = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix} = \Phi(\phi/4) = e^{i\pi/8}\text{R}_z(\phi/4) \quad (4.48)$$

We explicitly point out that, as it is easily understood from the definition in Eq. (4.41), any $\text{SF}_{N,j}$ is the inverse of itself. The full sequence realizing U_i can then be summarized as follows: starting from the initialized register $|0\rangle^{\otimes N}$, parallel Hadamard (H) gates are applied to create an equal superposition of all the elements of the computational basis

$$|0\rangle^{\otimes N} \xrightarrow{\text{H}^{\otimes N}} \frac{1}{\sqrt{m}} \sum_{j=0}^{m-1} |j\rangle \equiv |+\rangle^{\otimes N} \quad (4.49)$$

where $\sqrt{2}|+\rangle = |0\rangle + |1\rangle$. Then, the $\text{SF}_{N,j}$ blocks are applied one by one whenever there is a -1 factor in front of $|j\rangle$ in the representation of the target $|\psi_i\rangle$. Notice that any $\text{SF}_{N,j}$ only affects a single element of the computational basis while leaving all the others untouched. Moreover, all $\text{SF}_{N,j}$ blocks commute with each other, so they may actually be performed in any order.

The unitary U_w encoding the weight vector can be designed along the same lines. Indeed, given a $|\psi_w\rangle$, we first apply the $\text{SF}_{N,j}$ blocks that would be needed to flip all the -1 signs in its expansion

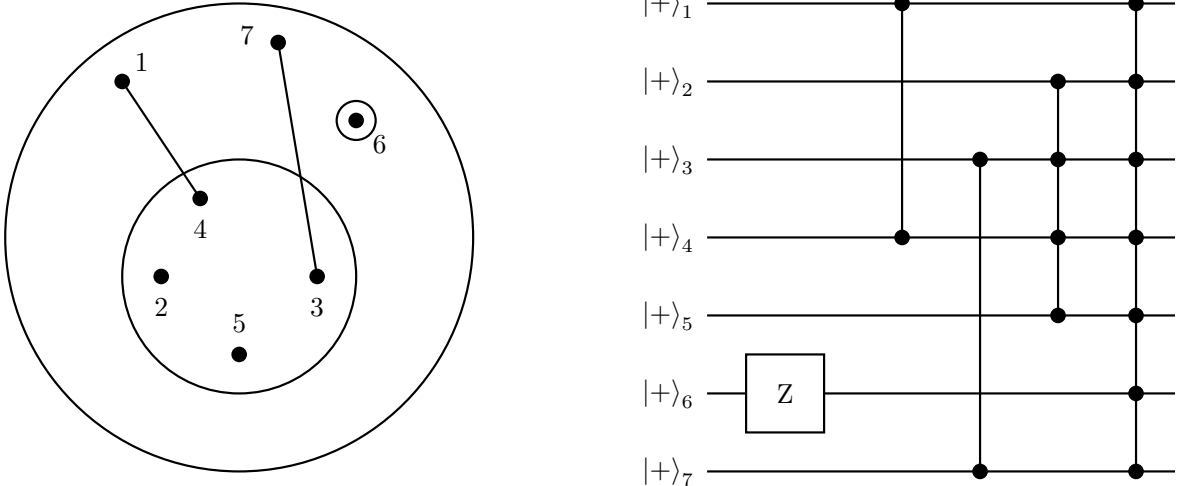


Figure 4.4: Quantum hypergraph states. A mathematical hypergraph, with hyper-edges represented as circles around subsets of vertices, can be put in correspondence with a quantum circuit generating the corresponding quantum hypergraph state. Notice that, since all the operations commute with each other, the sequence of gates can be performed in any order.

on the computational basis, thus returning to the balanced superposition $|\psi_w\rangle \rightarrow |+\rangle^{\otimes N}$. This quantum state can then be brought into the desired $|11\dots 11\rangle \equiv |1\rangle^{\otimes N}$ state by applying parallel Hadamard and NOT gates:

$$|+\rangle^{\otimes N} \xrightarrow{H^{\otimes N}} |0\rangle^{\otimes N} \xrightarrow{\text{NOT}^{\otimes N}} |1\rangle^{\otimes N} \quad (4.50)$$

As commented at the end of the previous section, the whole problem is manifestly symmetric under the addition of a global -1 factor (i.e. $|\psi_i\rangle$ and $-|\psi_i\rangle$ are fully equivalent, as it is always the case in quantum mechanics). Hence, there can be only at most $m/2 = 2^{N-1}$ independent -1 factors, and 2^{N-1} sign flip blocks are needed in the worst case: as a result, the above strategy is in general exponentially expensive in terms of circuit depth as a function of the number of qubits, and requires an exponential number of N -controlled gates. At the same time, in terms of the size of the classical input, the quantum algorithm scales linearly with the number of bits m .

Quantum hypergraph states

An optimized implementation of U_i and U_w can be given after realizing that the class of possible input- and weight-encoding states, Eq. (4.34), coincides with the set of the so called quantum *hypergraph* states. These are ubiquitous ingredients of many renown quantum algorithms and have been extensively studied and theoretically characterized [371, 372].

A k -uniform hypergraph $g_k = \{V, E\}$ is defined as a collection of n vertices V with a set of k -hyper-edges E , where each k -hyper-edge connects exactly k vertices and the usual notion of a connected graph is obtained for $k = 2$. A k -uniform quantum hypergraph state can be associated to g_k in the following way: after assigning a qubit in the $|+\rangle$ state to each vertex, for any k -hyper-edge a (multi)-controlled Z operation is performed between all the qubits $\{i_1, \dots, i_k\}$ connected by that hyper-edge. We thus obtain

$$|g_k\rangle = \prod_{\{i_1, \dots, i_k\} \in E} C^k Z_{i_1, \dots, i_k} |+\rangle^{\otimes N} \quad (4.51)$$

where, with a little abuse of notation, we assume $C^2 Z \equiv CZ$ and $C^1 Z \equiv Z = R_z(\pi)$. Similarly, we define a hypergraph $g_{\leq N} = \{V, E\}$ as a set of N vertices V with a set E of hyper-edges of any order k , not necessarily uniform. A quantum state of N qubits can be associated to any hypergraph $g_{\leq N}$

by generalizing the same procedure given above for the k -uniform case, namely by performing a C^kZ gate for each hyper-edge in $g_{\leq N}$

$$|g_{\leq N}\rangle = \prod_{k=1}^N \prod_{\{i_1, \dots, i_k\} \in E} C^k Z_{i_1, \dots, i_k} |+\rangle^{\otimes N} \quad (4.52)$$

It is immediately clear that k -uniform hypergraph states form a subset of the class of hypergraph states. Moreover, for N qubits there are exactly $\mathcal{N} = 2^{2^N - 1}$ different hypergraph states. In Fig. 4.4 we show an example of a hypergraph and of the quantum circuit generating its corresponding state. A more complete technical discussion on the properties of hypergraph states can be found in Refs. [371, 372], where alternative definitions using the stabilizer formalism are also given.

From a formal point of view, our proposed encoding scheme assigns a boolean value representing a single classical bit to each element of a N -qubit computational basis. This value is converted into a ± 1 factor, i.e. a 0 or π phase, which is then put in front of each component in a balanced quantum superposition of all the computational basis elements. The quantum states which result from all possible choices of the phase factors are known as Real Equally Weighted (REW), and can be given a mathematical formulation: for any boolean function $f : \{0, 1\}^N \rightarrow \{0, 1\}$, the corresponding REW state is

$$|f\rangle = \frac{1}{2^{N/2}} \sum_{x=0}^{2^N-1} (-1)^{f(x)} |x\rangle \quad (4.53)$$

It is straightforward to see that all states introduced in Eq. (4.34) can be interpreted as REW states. The relation between quantum hypergraph and input-/weight-encoding states that we claimed at the beginning of this section can then be put in precise terms as follows [371]:

The set G_{\pm}^N of N -qubit REW states and the set $G_{\leq N} = \{|g_{\leq N}\rangle\}$ of quantum hypergraph states coincide.

Indeed, $G_{\leq N} \subseteq G_{\pm}$ since, as shown in Eq. (4.52), any $|g_{\leq N}\rangle$ is obtained from the REW state $|+\rangle^{\otimes N}$ by applying (multi)-controlled- Z operations, namely by introducing additional π phases in front some computational basis components. The opposite inclusion can be proven constructively, as we will show in a moment as part of the implementation of the target U_i and U_w unitary operations.

Encoding with hypergraph states

Leveraging the connection with quantum hypergraph states, we will now provide a strategy to realize the unitary U_i with at most a single N -controlled $C^N Z$ and a number of p -controlled $C^p Z$ gates with $p < N$.

We assume again that the quantum register starts in the blank state $|0\rangle^{\otimes N}$ and we apply an initial $H^{\otimes N}$ gate, see Eq. (4.49), to obtain the simplest REW state $|+\rangle^{\otimes N}$. Then, the algorithmic generation of hypergraph states takes a series of iterative steps to which we will collectively refer to as the ‘‘hypergraph states generation subroutine’’ (HSGS). First, we check, in the representation of $|\psi_i\rangle$ on the computational basis, whether there is any component with only one qubit in state $|1\rangle$ (i.e. of the form $|0 \dots 010 \dots 0\rangle$) requiring a -1 factor. If so, the corresponding single qubit Z gate is applied by targeting the only qubit in state $|1\rangle$. Notice that this might introduce additional -1 factors in front of states with more than one qubit in state $|1\rangle$. Then, for $p = 2, \dots, N$, we consider the components of the computational basis with exactly p qubits in state $|1\rangle$. For each of them, an additional -1 sign is introduced in front of its current amplitude (if it is needed and it was not previously introduced) by applying the corresponding $C^p Z$ between the p qubits in state $|1\rangle$. Similarly, if an unwanted sign is already present due to a previous step, this can be easily removed by applying the same $C^p Z$. Since $C^p Z$ acts non trivially only on the manifold with p or more qubits being in state $|1\rangle$, the signs of all the elements with a lower number of $|1\rangle$ components are left unchanged. By construction, when $p = N$ all the signs are the desired ones. Notice that, by associating to each $C^p Z$ the corresponding hyperedge as shown in Fig. 4.4, this also proves that




Figure 4.5: Example of HSGS-based quantum circuit for a $N = 4$ artificial neuron. The input vector has elements $i_0 = i_1 = -1$, and $i_j = 1$ for $j = 2, \dots, 15$, while the weight vector has elements $w_2 = w_3 = w_4 = -1$, and 1 in all other entries (remember that for $N = 4$ the artificial neuron is capable of processing $m = 2^4 = 16$ -dimensional input and weight vectors). The HSGS is realized inside the U_i block after the initial $H^{\otimes N}$ gate, and in the U_w block before the final $H^{\otimes N}$ and $\text{NOT}^{\otimes N}$ operations.

each and every one REW state has a corresponding hypergraph, thus completing the equivalence $G_{\pm}^N = G_{\leq N}$.

As in the sign flip case, U_w can be obtained by slightly modifying the sequence of gates that would be used to generate $|\psi_w\rangle$. Indeed, one can start by first performing the HSGS tailored according to the ± 1 factors in $|\psi_w\rangle$: since all the gates involved in HSGS are the inverse of themselves and commute with each other, this step is equivalent to the unitary transformation bringing $|\psi_w\rangle$ back to the equally balanced superposition of the computational basis states $|+\rangle^{\otimes N}$. The desired transformation U_w is finally completed by adding parallel $H^{\otimes N}$ and $\text{NOT}^{\otimes N}$ gates, see Eq. (4.50).

An example of the full sequence for a specific $N = 4$ case is shown in Fig. 4.5. The HSGS-based algorithm successfully reduces the required quantum resources with respect to the brute force approach based on sign flip blocks: however, it may still involve an exponential cost in terms of circuit depth, i.e. of temporal steps of the quantum circuit when all possible parallelization of independent operations on the qubits is taken into account. Indeed, the sign flip algorithm described above requires $O(2^N)$ N -controlled Z gates when running on N qubits, in the worst case. Since any $C^N Z$ can be decomposed into $\text{poly}(N)$ elementary single and two-qubit gates [5], the overall scaling of the sign flip approach is $O(\text{poly}(N)2^N)$. On the other hand, the worst case for the HSGS, namely the fully connected hypergraph with N vertices, corresponds to applying once all the possible Z and $C^p Z$ operations for $2 \leq p \leq N$. Since all these operations commute, they can be arranged in such a way that all the qubits are always involved either in a single-qubit operation or a multi-controlled one (e.g., a Z on a certain qubit and the $C^{N-1} Z$ on the remaining ones can be done in parallel), for any progressive clock cycle. The overall number of clock cycles is still $O(2^N)$, as in the previous case, but now at most one slice contains a N -qubit operation, while all other slices with $p < N$ can be decomposed into $\text{poly}(p)$ elementary operations. In this respect, the proposed HSGS optimizes the number of multi-qubit operations and may result in a significant practical advantage in currently available superconducting quantum processors, where multi-qubit operations are not natively available.

Before proceeding, it is finally worth pointing out that the role of U_w can be interpreted *a posteriori* as that of canceling some of the transformations performed to prepare $|\psi_i\rangle$, or even all of them if the condition $\vec{i} = \vec{w}$ is satisfied. Further optimization of the algorithm, lying beyond the scope of the present work, might therefore be pursued at the compiling stage. However, we notice that the input and weight vectors can, in practical applications, remain unknown or hidden until runtime.

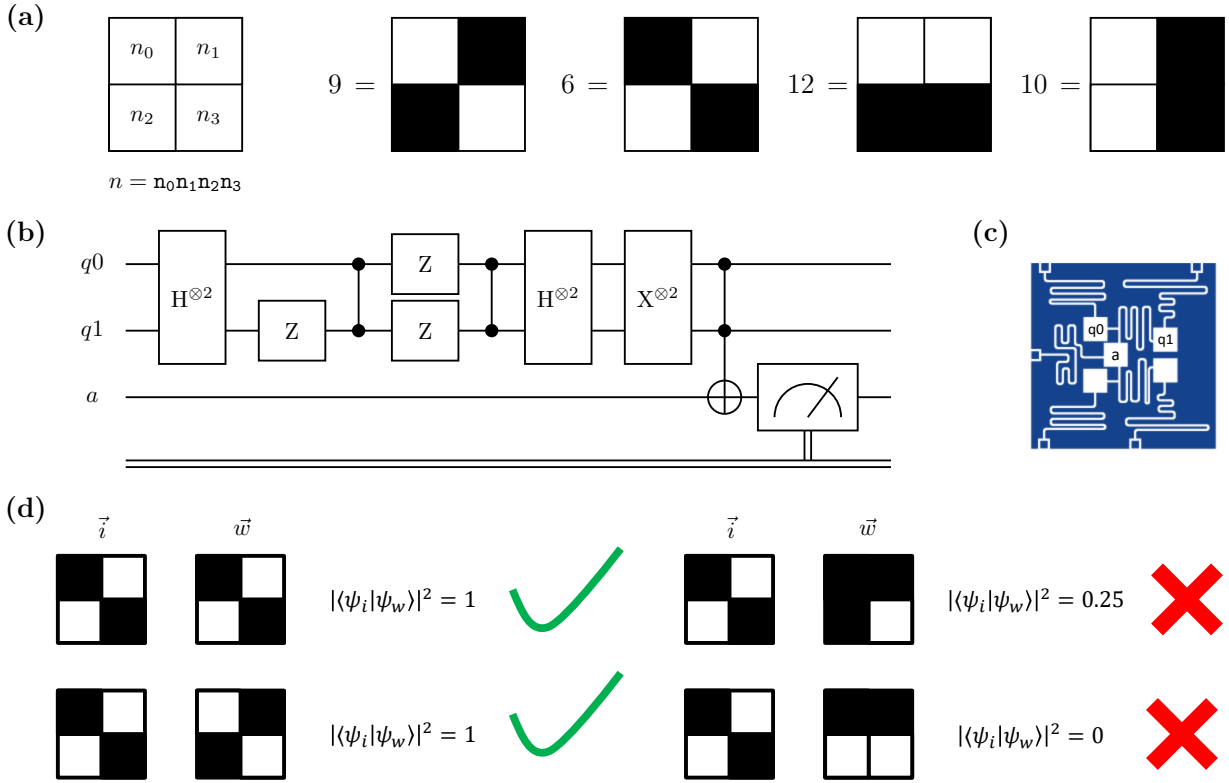


Figure 4.6: Pattern recognition with a single quantum artificial neuron. (a) Labeling of the black and white 2×2 visual patterns, with a few examples. (b) Example of the full gate sequence for the $N = 2$ case, with input and weight vectors corresponding to labels $k_i = 11$ and $k_w = 7$. (c) Mapping on the ibmqx4 “Tenerife” quantum processor. (d) Ideal recognition of a checkerboard pattern and its negative against other possible inputs.

4.2.2 Proof-of-principle realization on NISQ processors

We implemented the algorithm for a single quantum artificial neuron both on classical simulators working out the matrix algebra of the circuit and on the cloud-based real quantum backends provided via the IBM Quantum Experience (see Sec. 2.1). Due to the constraints imposed by the actual IBM hardware in terms of the size of the processors and the connectivity between the different qubits, we limited the quantum realization to the $N = 2$ case: nevertheless, even this small-scale example is already sufficient to show all the distinctive features of our proposed set up, such as the exponential growth of the analyzable problems dimension, as well as the pattern recognition potential. In general, as already mentioned, in this encoding scheme N qubits can store and process 2^N -dimensional input and weight vectors. Thus, 2^{2^N} different input patterns can be analyzed, corresponding also to the number of different \vec{w} that can in principle be defined, though the global phase symmetry in practice reduces to $\mathcal{N} = 2^{2^N - 1}$ the number of independent inputs. Given their essentially boolean nature, all binary inputs and weights can also easily be converted into black and white patterns, thus providing a visual interpretation of the artificial neuron classification and recognition activity.

When $N = 2$, $2^2 = 4$ binary images can be managed, and thus $2^{2^2} = 16$ different patterns can be analyzed. The conversion between \vec{i} or \vec{w} and 2×2 pixels visual patterns is done as follows: as depicted in Fig. 4.6a, we label each image by ordering the pixels left to right, top to bottom, and assigning a value $n_j = 1(0)$ to a white (black) pixel. The corresponding input or weight vector is then built by setting $i_j = (-1)^{n_j}$, or $w_j = (-1)^{n_j}$. We can also univocally assign an integer label k_i (or k_w) to any pattern by converting the binary string $n_0 n_1 n_2 n_3$ to its corresponding decimal number representation. Under this scheme, numbers 3 and 12 are for example used to label patterns with horizontal lines, while 5 and 10 denote patterns with vertical lines and 6 and 9 are used to label




Figure 4.7: Experimental results for the $N = 2$ case. (a) Ideal outcome of the quantum algorithm for an artificial neuron simulated on a classical computer. (b) Results from the ibmqx4 quantum processor using the 3-qubit algorithm with sign flip blocks. (c) Results from the ibmqx4 quantum processor using the 3-qubit HSGS-based algorithm. The average discrepancies \mathcal{D} calculated with respect to the ideal case are explicitly reported.

images with checkerboard-like pattern. An example of the sequence of operations performed on the IBM Q quantum computer using hypergraph states is shown in Fig. 4.6b for \vec{i} corresponding to the index $k_i = 11$, and \vec{w} corresponding to $k_w = 7$. The quantum algorithm was run on three qubits (including the ancilla) and all the gate decomposition techniques presented e.g. in Sec. 4.2.1 were used to explicitly reduce the quantum unitary transformations to the set of native IBM Q operations. Moreover, outcome probabilities for the final measure on the ancilla were reconstructed with $n_{\text{shots}} = 8192$ repetitions.

The Hilbert space of two qubits is relatively small, with a total of 16 possible values for \vec{i} and \vec{w} . Hence, the quantum artificial neuron model could be experimentally tested on the IBM quantum computer for all possible combinations of input and weights. The ideal outcome of the quantum algorithm is shown in Fig. 4.7a, where both the global -1 factor¹ and the input-weight symmetries are immediately evident. In particular, it can be realized that, for any given weight vector \vec{w} , the artificial neuron is able to single out from all the 16 possible input patterns only $\vec{i} = \vec{w}$ and its negative (with output $|c_{m-1}|^2 = 1$, i.e. the perfect activation of the neuron), while all other inputs give outputs (i.e. activation probabilities) smaller than 0.25. If the inputs and weights are translated into 2×2 black and white pixel grids thorough the set of rules adopted above, it is not difficult to see that a single quantum neuron can be used to recognize either vertical, horizontal lines, or checkerboard patterns (an example is reported in Fig. 4.6d). Notice however that the single node can only focus on one specific pattern at a time: for example, it cannot tell apart patterns containing *either* vertical or horizontal lines from all the others. The latter is a task that requires a deeper network, and will be tackled in Sec. 4.3.

The actual experimental results, obtained on the ibmqx4 “Tenerife” quantum processor, are shown in Fig. 4.7b-c. First (panel 4.7b), we report the results of the brute-force sign-flip. To demonstrate explicitly the progressive degrading of performances with increasing circuit length, here we deliberately did not take into account the global sign symmetry, thus treating any $|\psi_i\rangle$ and $-|\psi_i\rangle$ as distinct input quantum states and using up to 2^N sign flip blocks: even in such an elementary example the algorithm performs worse and worse with increasing number of blocks. It should however be emphasized that, despite the quantitative inaccuracy, the underlying structure of the output is already quite clear: indeed, a threshold of 0.5 applied to the measured output would be sufficient to successfully complete all the classification tasks, i.e. the artificial neuron can correctly single out from all possible inputs any given precalculated weight vector. A remarkably better quantitative accuracy is obtained using the algorithm based on the hypergraph states formalism,

¹With the integer labeling defined above, the global phase symmetry gets reflected in the fact that the outcome is symmetric for complementary labels modulo $2^{2^N} - 1 = 15$.

see panel 4.7c. In this case, the global phase symmetry is naturally embedded in the algorithm itself, and the results show symmetric performances all over the range of possible inputs and weights. All combinations of \vec{i} and \vec{w} yield results either larger than 0.75 or smaller than 0.3, in good agreement with the expected values, and all the classification tasks are once again correctly carried out. To give a quantitative measure of the overall agreement between the ideal (Fig. 4.7a), sign flips (Fig. 4.7b) and hypergraph states (Fig. 4.7c) versions, we introduce the average discrepancy

$$\mathcal{D} = \frac{\sum_{i,w} |O(i,w) - O_{ideal}(i,w)|}{2^{2^{N+1}}} \quad (4.54)$$

where $O(i,w) = \left| \sum_j i_j w_j \right|^2 = |c_{m-1}|^2$ is the outcome of the artificial neuron for a given pair of input and weights as obtained on a real device, $O_{ideal}(i,w)$ is the corresponding ideal result and $2^{2^{N+1}}$ is the total number of possible \vec{i} - \vec{w} pairs. As reported also in Fig. 4.7b-c, we find $\mathcal{D} \simeq 0.2364$ for the sign flip case, and $\mathcal{D} \simeq 0.0598$ for the HSGS-based version, thus quantitatively confirming the visual impression that the latter leads to better performances.

4.2.3 Hybrid training scheme

In the spirit of showing the potential scalability and usefulness of our quantum artificial neuron model, we have applied the HSGS-based algorithm to the $N = 4$ case by using the circuit simulator feature available in Qiskit [196]. Now 2^{32} possible combinations of \vec{i} and \vec{w} vectors are possible: to explicitly show a few examples, we have chosen a single weight vector, \vec{w}_t , corresponding to a simple cross-shaped pattern when represented as a 4×4 pixels image encoded along the same lines of the $N = 2$ case, and weighted it against several possible choices of input vectors. When a threshold $O(i,w_t) > 0.5$ is applied to the outcome of the artificial neuron, 274 over the total 2^{16} possible inputs are selected as positive cases, and they correspond to patterns differing from \vec{w}_t (or from its complementary one) by at most two pixels. Geometrically speaking, the vectors corresponding to positive cases all lie within a cone around \vec{w}_t . Some results are reported in Fig. 4.8a for a selected choice of input vectors, where the artificial neuron output is computed both with standard linear algebra and with a quantum circuit on a virtual and noise-free quantum simulator run on a classical computer. The discrepancies between the two are solely due to the finite statistics of the final measure required in the second case.

The true power of artificial neurons lies of course in their trainability. Based on the results above, we have taken a first step in this direction by implementing an elementary hybrid quantum-classical training scheme, through an adaptation of the perceptron update rule [373] to our algorithm. After preparing a random training set containing a total of 3050 different inputs, of which 50 positive and 3000 negative ones according to the threshold defined in the previous paragraph, the binary valued artificial neuron is trained to recognize the targeted \vec{w}_t . This is obtained by using the noiseless Qiskit simulator feature, in which the artificial neuron output is computed through our proposed quantum algorithm, and the optimization of the weight vector is performed by a classical processor. We selected a random \vec{w}_0 vector to start with, and then we let the artificial neuron process the training set according to well defined rules and learning rates l_p and l_n for positive and negative cases, respectively, without ever conveying explicit information about the target \vec{w}_t .

More in detail, at any step of an online² training session the current weight vector of the artificial neuron (\vec{w}) is modified or left unchanged depending on the answers provided by the artificial neuron. Whenever the outcome for a given \vec{i} in the training set and the current \vec{w} , $O(i,w)$, is correctly above or below the threshold fixed at the beginning to discriminate between positive and negative cases, \vec{w} is left unchanged. On the contrary, \vec{w} is updated when the input is wrongly classified: in particular, if the artificial neuron declares an input \vec{i} as a positive case ($O(i,w) > 0.5$ here) while it is actually labeled as a negative case in the precalculated training set, \vec{w} is moved further apart

²With *online* training, we mean that the elements in the training set are streamed in a sequential fashion. The weight vector is updated after every input has been processed and the actual positive/negative decision based on the outcome of the artificial neuron has been compared with the ideal one assigned to \vec{i} in the training set.




Figure 4.8: Training a $N = 4$ artificial neuron. (a) A possible cross-shaped choice for the target weight vector \vec{w}_t is represented together with a small selection of different input vectors. Above each input pattern, the quantitative answers of the artificial neuron, namely the values of $|c_{m-1}|^2$, are reported as obtained either through standard linear algebra (ideal ‘exact’ results) or resulting from the simulation of the quantum algorithm (‘q. alg’, run on a classical computer and averaged over $n_{\text{shots}} = 8192$ repetitions). The two versions agree within statistical error. The inputs with outcome $O(i, w) > 0.5$ differ from \vec{w}_t or its complementary by at most two pixels and are labeled as positive cases in the construction of the training set. (b) Example of a learning trajectory from the initial pattern \vec{w}_0 , reaching the target \vec{w}_t through some intermediate steps. (c) Average fidelity of the quantum state encoding the learned pattern with respect to the target one, obtained by repeating the learning procedure 500 times on the same training set.

from \vec{i} by flipping a fraction l_n of the ± 1 signs, randomly selected among those positions where the components of \vec{i} and \vec{w} coincide. Similarly, whenever a positive input \vec{i} is wrongly recognized as a negative one, \vec{w} is moved closer to \vec{i} by flipping a fraction l_p of the ± 1 signs in \vec{w} randomly chosen among those positions where the components of \vec{i} and \vec{w} differ from each other. The classical part of the training is performed within a single geometrical m -dimensional hemisphere, thus avoiding confusion between a pattern and its negative, although notice that this distinction is unnecessary in specifying the parameters of the quantum part of the algorithm. An example of the trajectory of the system around the configuration space of possible patterns is shown in Fig. 4.8b, in which we computed the fidelity of the quantum state $|\psi_w\rangle$ encoding the trained \vec{w} with respect to the target state $|\psi_{w_t}\rangle$. We point out explicitly that such fidelity is offered here as a convenient mean to depict the convergence of the training method, but that it is never used as a figure of merit during the training procedure: no information about the target \vec{w}_t other than the positive/negative labels assigned to elements in the training set is ever conveyed to the artificial neuron. In Fig. 4.8c, we also report the average behavior over 500 realization of the training scheme, all with the same initial pattern \vec{w}_0 and the same training set. As it can be seen, the quantum artificial neuron effectively learns the targeted cross-shaped pattern. With some heuristic adjustments, we find $l_p = 0.5$ to be the optimal learning rate in our case, with little effect produced by l_n . The latter is also set to $l_n = 0.5$ for simplicity.

4.3 A quantum feed-forward neural network

Moving from the simple, single layer artificial neuron design to complete feed-forward networks truly revolutionized classical machine learning. In feed-forward Neural Networks (ffNNs) there generally exists at least one additional layer of neurons (i.e. one or more “hidden” layers) between the input and the output ones, as represented in Fig. 4.2b. The advantage of ffNNs with respect to simpler designs such as single layer perceptrons or support vector machines is that they can be used to classify data with relations that cannot be reduced to a separating hyperplane. More formally, ffNNs can realize large classes of functions of the inputs under fairly general assumptions on the activation function of single nodes [333, 334].

In order to harness the full potentialities that quantum computing may offer to the field of artificial intelligence, it may well be necessary to undergo the same passage from single layered to deep feed-forward neural networks. After introducing in the previous section a possible, NISQ-accessible model of single artificial neurons capable of analyzing exponential amounts of classical data, here we take a first step towards a fully quantum coherent ffNN and we report some experimental tests on a state-of-the art 20-qubit IBM Q quantum processor. We use the HSGS-powered algorithm described in Sec. 4.2.1, combined with suitable quantum operations mimicking “synapses”, to build the network architecture. Although the procedure can in principle be extended to arbitrary network design, we will limit our analysis to a specific example whose structure already contains many characteristic features of ffNNs while remaining sufficiently simple to be successfully simulated on the present generation of superconducting quantum hardware.

A scheme of the ffNN hereby considered is shown in Fig. 4.9a, where the circles indicate artificial neurons, and the vectors \vec{w}_i refer to the respective weights, which we impose from the beginning. The network is constructed with a single hidden layer made of $N = 2$ neurons and a single binary (i.e. yes/no) output neuron, and is meant to carry out an elementary but meaningful problem: the network should recognize (i.e. give a positive answer) whether there exist straight lines in 2×2 images containing only black or white square pixels, regardless of the fact that the lines are horizontal or vertical. All the other possible input images should be classified as negative. As we saw above in Sec. 4.2.2, such images can indeed be encoded and processed by our quantum model of artificial neurons, but a single neuron cannot recognize at the same time vertical and horizontal lines: indeed, such patterns are, both literally and in vector representation, orthogonal to each other, and for each type there are two possible images representing the negative of each other. As a result, there is no single hyperplane separating just the four positive states from all other possible inputs.

The working principle of the network is straightforward. With the choice of weights represented in Fig. 4.9a, the top quantum neuron of the hidden layer outputs a high activation probability if the input vector has vertical lines, while the bottom neuron does the same for the case of horizontal lines. The output neuron in the last layer then combines the two pieces of information deciding whether one of the neurons in the hidden layer has given a positive outcome. Featuring a set of precalculated weights, our experiment is meant as a proof-of-principle demonstration of the fact that ffNNs can be run successfully on quantum processors. A more realistic use case, which is left for future work, would of course involve the full training of the network, for which a meaningful procedure should be designed and the network functionality probably extended to allow e.g. for continuous-valued data.

4.3.1 Connecting quantum artificial neurons

When several copies of the quantum register implementing the quantum artificial neuron model outlined in Sec. 4.2 work in parallel, the respective ancillae can be used to feed-forward the information about the input-weight processing to a successive layer, thus constructing a fully coherent quantum ffNN. By measuring the output layer only, the activation state of the network can be assessed, and several computational tasks that are in principle inaccessible to a single neuron can be performed.

A synapse, namely a connection between a neuron n_a belonging to a given layer of the network and another neuron n_b belonging to the successive layer, can be implemented as a controlled operation such that the state of the quantum register corresponding to n_b is conditioned upon the state




Figure 4.9: Architecture of a feed-forward neural network and corresponding quantum algorithm. (a) Model of a fFNN with four classical binary inputs, one hidden layer containing two artificial neurons and one output layer made of a single neuron. Next to each neuron, the ideal shape of the weight vector needed to achieve the desired recognition of horizontal and vertical lines is shown. The corresponding encoding scheme in terms of black and white pixels, which is the one introduced in Sec. 4.2.2, is also reported for convenience. (b) Quantum circuit implementing the fFNN. The first (second) three qubits encode the first (second) hidden neuron, while the last qubit represents the output layer. Synapses between the successive layers are color-coded, in agreement with the panel above. (c) Layout of the 20-qubit IBM Q Poughkeepsie processor, with qubit-qubit CNOT connectivity shown explicitly. The qubits used to run the fFNN quantum algorithm are highlighted and color-coded.

of the ancillae carrying the information about the activation state of n_a . An example of a practical implementation of this proposed architecture is given in Fig. 4.9b. We assume that each neuron within the hidden layer can accept 4-dimensional inputs, such that each quantum neuron can be represented on a 2-qubit encoding register plus an ancilla qubit (i.e. $m = 4$ and $N = 2$ using the notation of Sec. 4.2.1). At the same time, the output neuron takes 2-dimensional inputs coming from the previous layer and provides the global activation state of the network, thus requiring a single qubit ($m = 2$, $N = 1$) to be encoded. Let us call n_1 and n_2 the first two neurons and n_3 the output one, and assume that after the single-neuron quantum operations the global state of the $(3+3+1)$ -qubit network can be written as

$$\begin{aligned}
 & (r_{n_1} |\varphi_{n_1}\rangle |0\rangle_{a_1} + c_{m-1, n_1} |1 \dots 1\rangle_{n_1} |1\rangle_{a_1}) \\
 & \otimes (r_{n_2} |\varphi_{n_2}\rangle |0\rangle_{a_2} + c_{m-1, n_2} |1 \dots 1\rangle_{n_2} |1\rangle_{a_2}) \\
 & \otimes |0\rangle_{n_3}
 \end{aligned} \tag{4.55}$$

where we explicitly indicated with a_x the ancilla for neuron n_x , $r_{n_x} = (1 - c_{m-1, n_x}^2)^{1/2}$ and $|\varphi_{n_x}\rangle$ contains, for each neuron, all the components other than the one leading to activation, see Eq. (4.40). Notice that, by construction, $\langle \varphi_{n_x} | 1 \dots 1 \rangle = 0$. In the meantime, the n_3 qubit is brought into the superposition $\sqrt{2}|+\rangle = |0\rangle + |1\rangle$ by applying a single-qubit Hadamard gate, H. Synapses can thereafter be implemented with two CZ gates, as represented in Fig. 4.9b. The overall state of the quantum fNN then becomes

$$\begin{aligned} & (r_{n_1} r_{n_2} |R_{n_1}\rangle |R_{n_2}\rangle + c_{n_1} c_{n_2} |A_{n_1}\rangle |A_{n_2}\rangle) |+\rangle_{n_3} \\ & + (r_{n_1} c_{n_2} |R_{n_1}\rangle |A_{n_2}\rangle + c_{n_1} r_{n_2} |A_{n_1}\rangle |R_{n_2}\rangle) |-\rangle_{n_3} \end{aligned} \quad (4.56)$$

where c_{n_x} is a short-hand notation for c_{m-1, n_x} , and the activated $|A\rangle$ and rest $|R\rangle$ states of n_1 and n_2 are explicitly given as

$$\begin{aligned} |A_{n_x}\rangle &= |1 \dots 1\rangle_{n_x} |1\rangle_{a_x} \\ |R_{n_x}\rangle &= |\varphi_{n_x}\rangle |0\rangle_{a_x} \end{aligned} \quad (4.57)$$

It is now worth pointing out the following aspects of this original quantum fNN formulation. First, a general superposition of all the possible independent encoding REW states with $m = 2$ is constructed for the output neuron, conditioned on the state of the quantum registers in the previous layer, through the control synapses. Thus, one can now proceed to the analysis of the inputs on n_3 through a unitary operation U_{w_3} , according to the general algorithm already outlined for single quantum neurons. Moreover, as it will be explicitly shown below in a practical example, in this case the output neuron can also be measured directly without adding further ancillae qubits. Second, different combinations of activation and rest patterns leading to a given state of the output neuron can be engineered by slightly modifying the structure of the synapses: as an example, if the n_1 - n_3 connections were given by a generalized CZ operation, controlled by a_1 being in state $|0\rangle$ (e.g. by adding a NOT gate on a_1 before and after the CZ [5]), then the final state of the quantum fNN would read

$$\begin{aligned} & (r_{n_1} r_{n_2} |R_{n_1}\rangle |R_{n_2}\rangle + c_{n_1} c_{n_2} |A_{n_1}\rangle |A_{n_2}\rangle) |-\rangle_{n_3} \\ & + (r_{n_1} c_{n_2} |R_{n_1}\rangle |A_{n_2}\rangle + c_{n_1} r_{n_2} |A_{n_1}\rangle |R_{n_2}\rangle) |+\rangle_{n_3} \end{aligned} \quad (4.58)$$

We notice, however, that the initial choice leading to Eq. (4.56) seems the most natural one, especially if the aim is to keep the classical interpretation of a synapse firing when the feeding neuron is found in its active state.

4.3.2 Pattern classification on a real quantum processor

Leveraging the remarkable performances of the IBM Q Poughkeepsie quantum processor, we realized in practice the 7-qubit experiment introduced in Fig. 4.9b. This is, to the best of our knowledge, one of the largest quantum neural network computation reported to date on real quantum hardware in terms of the total size of the quantum register.

All individual nodes of the network were treated with the HSGS-based procedure described in Sec. 4.2.1. In particular, the two neurons representing the hidden layer can work in parallel on two separate 3-qubit registers, which then need to separately interact with the single qubit encoding neuron n_3 . Given the still manageable size of the computation, all the mapping and connectivity adaptation of the circuit was hard coded without any automated compilation. In particular, here we notice that, at difference with the ibmqx4 triangular topology of Sec. 4.2.2, the two 3-qubit subsets encoding for a single node are only linearly connected. This makes the final Toffoli gate writing the activation state of each neuron on the respective ancilla not directly implementable via e.g. the decomposition reported in Fig. 4.3b, for that requires at least two CNOT operations between two qubits which are not nearest neighbors on the chip. We thus resort to SWAP constructions, via the identity

$$\begin{array}{c} \times \\ | \\ \times \end{array} = \begin{array}{c} \bullet \oplus \bullet \\ | \quad | \\ \oplus \bullet \oplus \end{array} \quad (4.59)$$

where the numbers represent physical qubits as depicted in Fig. 4.9c. The advantage in total circuit depth with respect to the exact construction of Fig. 4.3b, when SWAP operations are taken into account, is significant for the performances of a NISQ processor. An example of the full sequence of quantum gates run via the IBM Quantum Experience is reported in Fig. 4.10a. As shown in Fig. 4.9a, we choose \vec{w}_1 and \vec{w}_2 corresponding to a vertical and a horizontal line, respectively, and we set $\vec{w}_3 = (1, -1)$. As a result, the HSGS-based algorithm for quantum neurons gives $U_{w_3} \equiv H$ and $c_{n_1} = 1$ ($c_{n_2} = 1$) if and only if \vec{i} contains a vertical (horizontal) line. Going back to Eq. (4.56) and applying U_{w_3} on the quantum state of the register after the synapses, we obtain an output state

$$\begin{aligned} |\psi_{out}\rangle = & \\ & (r_{n_1} r_{n_2} |R_{n_1}\rangle |R_{n_2}\rangle + c_{n_1} c_{n_2} |A_{n_1}\rangle |A_{n_2}\rangle) |0\rangle_{n_3} \\ & + (r_{n_1} c_{n_2} |R_{n_1}\rangle |A_{n_2}\rangle + c_{n_1} r_{n_2} |A_{n_1}\rangle |R_{n_2}\rangle) |1\rangle_{n_3} \end{aligned} \quad (4.61)$$

Since the output neuron is encoded in a single qubit, we can directly read out its activation state by measuring it in the computational basis. The probability of finding $|1\rangle$ on n_3 is therefore

$$\begin{aligned} p_1 = & \text{Tr}[(\mathbb{I} \otimes |1\rangle\langle 1|_{n_3}) |\psi_{out}\rangle\langle\psi_{out}|] \\ = & r_{n_1}^2 c_{n_2}^2 + c_{n_1}^2 r_{n_2}^2 \end{aligned} \quad (4.62)$$

where \mathbb{I} is the identity operator on the quantum registers encoding neurons 1 and 2 with their respective ancillae. As shown in Fig. 4.10b, we ideally have $p_1 = 1$ for the target horizontal and vertical patterns, and $p_1 < 0.5$ in all other cases. In Fig. 4.10c we report the experimental results obtained from the IBM Q Poughkeepsie processor: as it can immediately be appreciated, despite some residual quantitative inaccuracy in the estimation of probabilities, all the classification tasks are correctly performed. In particular, the target horizontal and vertical patterns are singled out from all other possible inputs. Raw data from the quantum processor already allow for an accurate classification: however, the overall quality of the outcomes greatly benefits from the application of error mitigation techniques, as described in more detail in the following.

Error mitigation techniques

In Sec. 2.3.2 we have already introduced the concept of error mitigation as a way of improving the overall quality of experimental results obtained on NISQ processors. Here we describe how similar techniques can be applied in the present case to enhance the classification performances of our quantum fFNN algorithm and reduce the quantitative differences from the ideal outcome. In particular, here we cannot make use of general symmetry considerations or sum rules inspired by physical properties of the model under study, but we may apply some general-purpose protocols inspired by pulse-stretching and noise amplification methods described in the literature [62, 170, 172, 195].

We performed an incremental analysis on the performances of the circuit by artificially augmenting the effective level of noise affecting the computation. As end users on the IBM Quantum Experience, we did not have at the time we performed the experiments (June 2019) access nor control of the actual pulse scheme used to drive the qubits. We then resorted to the mathematical properties of e.g. CNOT gates, which are the inverse of themselves, to control the effective noise in discrete steps. Indeed, by selecting an odd integer $r = 1, 3, 5, \dots$ as a noise parameter, an increasingly noisy CNOT can be implemented by repeating the same gate r times:

$$\left(\begin{array}{c} \bullet \\ | \\ \oplus \end{array} \right)_{\text{noisy}} = \left(\begin{array}{c} \bullet \\ | \\ \oplus \end{array} \right)^r \quad (4.63)$$

By applying this procedure to all the CNOT gates in a given circuit (including those introduced to decompose e.g. CZ and SWAP gates) and for different values of r , we obtain a set of noisy outcomes as a function of r . Error mitigated data can then be constructed by extrapolating such results




Figure 4.11: Noise amplification and error mitigation (a) Results of the ffNN computation on IBM Q Pughkeepsie for noise amplification applied to CNOT gates and different choices of the noise parameters. (b) Results for noise amplification applied to CNOT and H gates. (c) Error mitigation results, extrapolated with a quadratic fit and compared with raw data for the case on which the noise amplification is applied to CNOT gates only or both CNOT and H operations. (d) Example of zero-noise extrapolation with polynomial fit of degree d for inputs with corresponding integer label i .

at zero noise. An example of the application of this technique to the data for the ffNN quantum computation is shown in Fig. 4.11a-c. The extrapolation is done assuming a quadratic fit function, see Fig. 4.11d.

The error mitigation scheme we have described is essentially focused on random noise affecting two-qubit operations. Errors on CNOT gates, particularly when the latter are implemented via the relatively slow cross-resonance procedure, are indeed a common source of inaccuracy and the procedure outlined above is therefore able to enhance the overall quality of the results obtained on the real quantum processor. However, we have seen and extensively motivated in Sec. 2.3.3 that single-qubit operations are also typically affected by significant amounts of errors whose impact on the results can be comparable, if not larger, to that of two-qubit gates. Although the most harmful effects in such case are actually coherent, and thus not easily tamed by noise amplification methods, over the typical duration of single-qubit pulses (up to 50-100 ns for Gaussian Derivative on IBM Q devices, see Sec. 2.1) incoherent quantum channels can have non-negligible impact. A further step can then be taken by combining the CNOT mitigation with a similar approach applied to H gates, which are ubiquitous in the proposed algorithm and have the self-inverse property $H^2 = I$. In Fig. 4.11b we report the results of the application of the r -fold noise incremental enhancement to all CNOT and H in the given circuit. The result of the extrapolation for both strategies (mitigation applied to CNOT alone, or to CNOT and H) are compared with raw results in Fig. 4.11c. The quality of the results for the CNOT+H case, which is the one presented in the main text, is improved for almost all data points with respect to the corresponding CNOT-only mitigation. Residual inaccuracies are most probably due to systematic biases introduced either in the realization

of individual coherent operations or during readout. We also notice that a more refined series of interpolation experiments [195] could in principle be used to estimate higher-order correction terms and to tackle non-stochastic errors.

Before concluding, it is worth mentioning that the fully coherent approach to ffNNs presented here can also be reformulated in a hybrid quantum-classical way, with single quantum nodes connected through classical feed-forward of information upon measurement of the activation states of individual neurons [307]. Such description makes more explicit the non-linear contribution brought about by each layer of the network and can be directly related to the fully coherent version via the well known principle of deferred measurement [5]. The latter states that in a quantum circuit one can always move a measurement done at an intermediate stage to the end of the computation while replacing classically controlled operations (i.e. classical feed-forward synapses in this case) with quantum controlled ones, which indeed appear e.g. in the quantum coherent formulation of Fig.4.9b. On one hand, a hybrid architecture immediately suggest a way towards integration with existing classical structures and algorithms for neural network computation and machine learning: for example, one could imagine that a few carefully distributed quantum nodes at the input of an otherwise classical network might act as a memory-efficient layer enabling the treatment of otherwise unmanageable sets of data. On the other hand, the quantum coherent version reduces the necessity to store and process classical bits during intermediate stages, and could more directly point towards extensions in non-classically accessible regimes either with respect to the size of the computation or the complexity of the convolutional filters to be deployed.

Conclusions

In this work, we presented the current status of digital quantum technologies for simulation and information processing, and we introduced a number of original algorithms, hardware solutions and applications. Here we briefly summarize the main results and we discuss possible future directions of research.

Progress in the field of digital quantum simulations, which was reviewed in Chapter 1 and within which the discussion in Chapters 2 and 3 can be framed, is certainly closely connected to the evolution of practical quantum computing at large. Indeed, we saw that a fully developed quantum architecture satisfying the requirements for the implementation of the standard circuit-based quantum computation can equally well represent a general-purpose quantum simulator, capable of reproducing efficiently the dynamics of a large class of physical systems. As they formally obey the algebra of qubits, many of the spin-1/2 models that we discussed and implemented on real NISQ processors may serve both as ideal test-beds for benchmarking specific hardware realizations and as the joining link towards more complex physical systems. A certified quantum advantage, nowadays the driving force behind a large number of actors in the quantum computing arena, would of course open in the long run exciting avenues in the physical sciences: in the language of simulations, it could for example mean approaching the unexplored frontier of highly correlated many-body systems. On a more near-term perspective, and provided that significant efforts are devoted to improve reliability and scalability of the current noisy prototypes, practical speed-ups may be expected in specific use-cases including quantum chemistry and mid-sized magnetic systems: in this direction, novel hybrid quantum-classical approaches, including variational quantum eigensolvers or time propagators and data post processing techniques such as the ones we described in Chapter 2 for the extraction of experimental cross sections, look particularly promising and sufficiently resilient. While exact time evolution, requiring long sequences of quantum gates and high degree of accuracy per digital step, will probably remain in the immediate future inaccessible for most non-trivial physical models, we have nevertheless indicated a possible path towards the use of prototypical quantum simulation methods for obtaining useful characteristic and structural quantities which are in principle hard to compute classically. It is also worth pointing out that crossing the line of what even the most powerful supercomputers are able to simulate will immediately pose the issue of verification. Once more, the world and its description will then need to be compared on equal footing: as long as physical systems are involved, experiments will ultimately certify the validity of computational results, although this could not always represent the most practical solution. A judicious characterization of the hardware working principles could certainly offer at least partial legitimacy, yet we will with great probability assist to an increasing interest into the development of formal validation and diagnostic tools. In the meanwhile, though on a much longer time-scale, the pursuit of fault tolerance will certainly continue.

Within the multitude of alternatives to trapped ions and superconducting circuits, the development of hybrid quantum hardware could offer viable ways to overcome the present limitations of quantum information processing architectures. Be it merely a matter of performances or a deeper

technological challenge, the versatility and interoperability of quantum devices certainly represent almost unexploited resources. While it is certainly fair to say that the current level of maturity of many perspective quantum technologies is not yet comparable with the leading ones, this in turn means that they might gain space and attention in academic and fundamental research while the most developed counterparts drift towards the industrial world. As an example, we designed in Chapter 3 a novel and intrinsically scalable electromechanical architecture with promising features in terms of key figures of merit. Our theoretical proposal is supported by extensive numerical investigations of the actual hardware functioning and by a steady advancement, well documented in the recent literature, of the experimental fabrication and control techniques for nano-oscillators integrated within superconducting platforms.

The most exploratory part of our work focused, in Chapter 4, on describing new applications of quantum computing beyond its traditional scope. In particular, we gave a brief account of the recent birth and fast growth of quantum machine learning (QML) as a new and promising discipline, with vast horizons and possibly far reaching consequences. We have proposed a quantum model for binary-valued artificial neural computation, which we have then experimentally tested both at the single node and feed-forward network level. Encouraging results were obtained on IBM Q real quantum devices with up to 7 active qubits. Our encoding scheme in principle guarantees an exponential memory advantage with respect to classical counterparts: however, we must also explicitly acknowledge that, in general, the computational cost for constructing the relevant unitary transformations and synapse connections is also exponential in the size of the quantum register, or linear in the number of bits of the classical input. This is indeed a quite common, and sometimes neglected, bottleneck in many QML proposals. In a NISQ perspective, the practical necessity to decompose complex quantum operations, such as the multiply controlled gates appearing in our procedure, into hardware-native equivalents may also represent a further obstacle, especially for those platforms such as superconducting quantum circuits which typically feature only pairwise interactions. Our algorithm is formally exact and platform-independent: as such, it could be interpreted as a reference standard. Taking inspiration from variational unitary constructions and approximate methods for state preparation, we may then envision in the future viable quasi-exact alternatives with better scaling properties. Furthermore, it is not redundant to point out that the way in which data are collected and processed might also undergo significant rethinking over time, for example with quantum memories and databases possibly becoming available and at least partially easing the demand for general quantum-classical interfaces. Finally, quantum approaches to learning are of course the most natural next step in the search for a fully quantum-enhanced neural computation: while we have demonstrated, via a hybrid procedure, that our model is trainable at the single neuron level, the real fundamental question on whether genuinely quantum resources such as superposition and parallelism could speed-up the deep learning process remains open. Applying to our quantum feed-forward network, or to a possible future analogue with graded-response nodes, a quantum version of the backpropagation algorithm would be a remarkable step in this direction.

All considered, our analysis depicted the lively and rich ecosystem flourishing around quantum information processing and its applications. While a number of formidable challenges still stand in front of us, we may say with confidence that quantum mechanics is no longer just a physical theory: instead, it nowadays represents the core of the technological quest that will define our times.

In this Appendix, we provide additional numerical results and experimental data to complete the discussion of Chapter 2.

A.1 Suzuki-Trotter decomposition

As discussed in Sec. 2.3.1, the Suzuki-Trotter formula was applied to approximate the time evolution induced by the spin-spin couplings and the effect of external fields. A number $n > 1$ of Trotter steps is required for molecules 2, 4 and 5 to get a reliable reproduction of the actual physical dynamics: as this is in practice limited by the performances of the hardware, we have adopted a fixed n approach, calibrating the digital error on the longest simulation times. When possible, different intervals requiring different number of digital steps were identified, in order to reach optimal compromise between accuracy and circuit depth. Although exact simulations may of course not be easily available for larger models, for our proof-of-principle demonstration we made sure through numerical tests that the digital error does not affect the physically relevant features, such as oscillation frequencies and beatings, in such a way that the actual capabilities of the hardware could be assessed without other induced biases. As shown in Fig. A.1 and A.2, the comparison between the exact system evolution and the digital evolution, obtained by Suzuki-Trotter decomposition, justifies the choices made in Sec. 2.3.1.




Figure A.1: Digital vs exact time evolution for molecule 2. Dynamical correlation functions $C_{ij}^{\alpha\beta}$ (real part): the correct dynamics is reproduced with $n = 2$ Trotter step for short times ($Jt \leq 2.0$ in our simulations) and with $n = 4$ for $2.0 < Jt \leq 6.0$. Here we used a second order Suzuki-Trotter decomposition.




Figure A.2: Digital vs exact time evolution for molecule 4. Dynamical correlation functions $C_{ij}^{\alpha\beta}$ for molecule 4 (real part). The correct dynamics is reproduced with $n = 1$ Trotter step for short times ($Jt \leq 2.0$). Conversely, a decomposition with $n = 2$, although not perfect, is able to catch the essential features of the system dynamics, reproducing the correct frequency and the beats in the oscillations.

A.2 Digital quantum simulations of dynamical correlations

In the following pages, we report additional experimental results obtained on IBM Q quantum processors to fully characterize the set of dynamical correlation functions for the spin models under study in Sec. 2.3. These results constitute the basis for the extraction of Fourier amplitudes and frequencies, namely of spin matrix elements and eigenenergies, and for their subsequent computation of the 4D-INS cross section. Although only a the most representative results are reported here, we mention that the total number of data points simulated via the IBM Quantum Experience is close to 10^4 .

All results shown here are already treated with the Phase-and-Scale error mitigation procedure introduced in Sec. 2.3.2. The data points are color coded with red representing real parts and blue representing imaginary parts of spin-spin dynamical correlation functions

$$C_{ij}^{\alpha\beta}(t) = \langle 0 | s_{\alpha}^{(i)}(t) s_{\beta}^{(j)} | 0 \rangle \quad (\text{A.1})$$

The exact results, obtained by diagonalizing the corresponding molecular spin Hamiltonian, are also reported for comparison.




Figure A.3: Dynamical correlation functions C_{ij}^{xx} for molecule 1 calculated on the ibmqx4 processor (circles), compared with the exact result (line).




Figure A.4: Dynamical correlation functions C_{ij}^{yy} and C_{ij}^{zz} for molecule 1 calculated on the ibmqx4 processor (circles), compared with the exact result (line).




Figure A.5: Dynamical correlation functions C_{ij}^{xy} for molecule 1 calculated on the `ibmqx4` processor (circles), compared with the exact result (line). Notice that contributions from the first and second line cancel out in the 4D-INS cross-section.




Figure A.6: Dynamical correlation functions C_{ij}^{zx} and C_{ij}^{zy} for molecule 1 calculated on the `ibmqx4` processor (circles), compared with the exact result (line).




Figure A.7: Dynamical correlation functions for molecule 1 calculated on the `ibmqx4` processor (circles), compared with the exact result (line). The experiment was extended to longer simulation times to check the periodic behavior of $C_{ij}^{\alpha\beta}(t)$.




Figure A.8: Dynamical correlation functions $C_{ij}^{xx}(t)$ for molecule 2 calculated on the `ibmqx4` processor (circles), compared with the exact result (line).




Figure A.9: Dynamical correlation functions $C_{ij}^{yy}(t)$ and $C_{ij}^{zz}(t)$ for molecule 2 calculated on the ibmqx4 processor (circles), compared with the exact result (line).




Figure A.10: Dynamical auto-correlation functions for molecule 3 calculated on the ibmqx4 processor (circles), compared with the exact result (line).




Figure A.11: Dynamical cross-correlation functions for molecule 3 calculated on the ibmqx4 processor (circles), compared with the exact result (line).




Figure A.12: Dynamical C_{ij}^{xy} and C_{ij}^{xz} correlation functions for molecule 3 calculated on the ibmqx4 processor (circles), compared with the exact result (line).




Figure A.13: Dynamical C_{ij}^{xx} correlation functions for molecule 4 (circles) calculated on the ibmqx5 (left column) and ibmqx4 (right column) processors, compared with the exact result (line).




Figure A.14: Dynamical C_{ij}^{yy} correlation functions for molecule 4 (circles) calculated on the ibmqx5 (left column) and ibmqx4 (right column) processors, compared with the exact result (line).




Figure A.15: Dynamical C_{ij}^{zz} correlation functions for molecule 4 (circles) calculated on the ibmqx5 (left column) and ibmqx4 (right column) processors, compared with the exact result (line).

A.3 Fitting dynamical correlations

Finally, we report in the following plots and tables some additional results obtained fitting the results of the digital quantum simulations of dynamical correlations on IBM Q processors, according to the Fourier expansion of Eq. (2.30)

$$C_{ij}^{\alpha\beta}(t) = \sum_p \langle 0 | s_\alpha^{(i)} | p \rangle \langle p | s_\beta^{(j)} | 0 \rangle e^{-iE_p t} \quad (\text{A.2})$$

As in the previous section, all experimental data are already error-mitigated via the PaS procedure.




Figure A.16: Fit of dynamical C_{22}^{xx} correlation function for molecule 2. Real (top) and imaginary (bottom) parts are fitted with a superposition of two oscillating functions (lines).




Figure A.17: Fit of dynamical auto-correlation functions C_{ii}^{xx} for molecule 3. Real (top) and imaginary (bottom) parts are fitted with a single oscillating function, indicating that the excitations of the two spins are fully independent.




Figure A.18: Fit of the imaginary part of dynamical C_{ij}^{xx} correlation functions for molecule 4 with a superposition of three sinusoidal functions, using the same amplitudes as for real parts reported in Fig. 2.7e.

Table A.1: Oscillation frequencies and Fourier coefficients of the dynamical correlations functions for the isotropic trimer (molecule 4), obtained by fitting the data measured on the ibmqx4 and ibmqx5 processors. Exact results obtained from spin Hamiltonian diagonalization are reported in square brackets. Values for symmetry-equivalent sites are not reported.

$ p\rangle$	E_p/J	$ \langle 0 s_x^{(1)} p\rangle ^2$	$ \langle 0 s_x^{(2)} p\rangle ^2$	$\langle 0 s_x^{(1)} p\rangle\langle p s_x^{(2)} 0\rangle$	$\langle 0 s_x^{(1)} p\rangle\langle p s_x^{(3)} 0\rangle$
$ 1\rangle$	8.5(1) [8.5]	0.04(3) [0.0425]	0.13(5) [0.1675]	-0.08(3) [-0.0825]	0.05(3) [0.0425]
$ 2\rangle$	9.5(1) [9.5]	0.14(2) [0.125]	0.02(3) [0.00]	-0.02(3) [0.00]	-0.13(1) [-0.125]
$ 3\rangle$	10.0(1) [10.0]	0.08(3) [0.0825]	0.11(4) [0.0825]	0.09(3) [0.0825]	0.08(3) [0.0825]

Table A.2: Oscillation frequencies and Fourier coefficients of the dynamical correlations functions for the anisotropic trimer (molecule 5), obtained by fitting the data measured on the ibmqx4 and ibmq20 processors with Eq. (2) of the main text. Exact results obtained from spin Hamiltonian diagonalization are reported in square brackets.

	$ p\rangle = 1\rangle$	$ p\rangle = 2\rangle$	$ p\rangle = 3\rangle$
E_p/J	8.0(1) [7.97]	9.1(1) [9.15]	9.5(1) [9.48]
$ \langle 0 s_x^{(1)} p\rangle ^2$	0.1(1) [0.18]	0.4(1) [0.63]	0.2(1) [0.19]
$ \langle 0 s_x^{(2)} p\rangle ^2$	0.5(1) [0.77]	0.0(1) [0.06]	0.2(1) [0.18]
$ \langle 0 s_x^{(3)} p\rangle ^2$	0.0(1) [0.05]	0.2(1) [0.31]	0.5(1) [0.64]
$\langle 0 s_x^{(1)} p\rangle\langle p s_x^{(2)} 0\rangle$	0.3(1) [0.37]	-0.2(1) [-0.19]	-0.2(1) [-0.18]
$\langle 0 s_x^{(2)} p\rangle\langle p s_x^{(3)} 0\rangle$	0.3(1) [0.20]	0.1(1) [0.13]	-0.2(1) [-0.33]
$\langle 0 s_x^{(1)} p\rangle\langle p s_x^{(3)} 0\rangle$	0.1(1) [0.10]	-0.4(1) [-0.44]	0.2(1) [0.35]

Bibliography

- [1] J. S. Bell, *Speakable and Unsayable in Quantum Mechanics*, Cambridge University Press, Cambridge, UK (1987).
- [2] S. Lloyd, *Computational Capacity of the Universe*, Phys. Rev. Lett. **88**, 237901 (2002).
- [3] R. Landauer, *Dissipation and noise immunity in computation and communication*, Nature **335**, 779–784 (1988).
- [4] R. Landauer, *The physical nature of information*, Phys. Lett. A **217**, 188–193 (1996).
- [5] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, Cambridge University Press, Cambridge, UK (2000).
- [6] J. Preskill, *Quantum Computing in the NISQ era and beyond*, Quantum **2**, 79 (2018).
- [7] J. Gubernatis, N. Kawashima and P. Werner, *Quantum Monte Carlo Methods*, Cambridge University Press, Cambridge, UK (2016).
- [8] J. M. Haile, *Molecular Dynamics Simulation: Elementary Methods*, Wiley-Interscience, New York, USA (1992).
- [9] S. Montangero, *Introduction to Tensor Network Methods*, Springer Nature Switzerland AG, Cham, CH (2018).
- [10] R. P. Feynman, *Simulating physics with computers*, Int. J. Theor. Phys. **21**, 467 (1982).
- [11] Y. I. Manin, *Computable and Noncomputable*, Sov. Radio p. 13 (1980).
- [12] P. Benioff, *The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines*, J. Stat. Phys. **22**, 563 (1980).
- [13] M. Fisher, P. B. Weichman, G. Grinstein *et al.*, *Boson localization and the superfluid-insulator transition*, Phys. Rev. B **40**, 546–570 (1989).
- [14] D. Jaksch, C. Bruder, J. I. Cirac *et al.*, *Cold Bosonic Atoms in Optical Lattices*, Phys. Rev. Lett. **81**, 3108–3111 (1998).
- [15] M. J. Hartmann, F. G. S. L. Brandao and M. B. Plenio, *Strongly interacting polaritons in coupled arrays of cavities*, Nat. Phys. **2**, 848 (2006).
- [16] A. D. Greentree, C. Tahan, J. H. Cole *et al.*, *Quantum phase transitions of light*, Nat. Phys. **2**, 856 (2006).

- [17] D. G. Angelakis, M. F. Santos and S. Bose, *Photon-blockade-induced Mott transitions and XY spin models in coupled cavity arrays*, Phys. Rev. A **76**, 031805 (2007).
- [18] I. Carusotto, S. Fagnocchi, A. Recati *et al.*, *Numerical observation of Hawking radiation from acoustic black holes in atomic Bose–Einstein condensates*, New J. Phys. **10**, 103001 (2008).
- [19] I. Carusotto, D. Gerace, H. E. Tureci *et al.*, *Fermionized Photons in an Array of Driven Dissipative Nonlinear Cavities*, Phys. Rev. Lett. **103**, 033601 (2009).
- [20] L. Tian, *Circuit QED and Sudden Phase Switching in a Superconducting Qubit Array*, Phys. Rev. Lett. **105**, 167001 (2010).
- [21] D. Gerace and I. Carusotto, *Analog Hawking radiation from an acoustic black hole in a flowing polariton superfluid*, Phys. Rev. B **86**, 144505 (2012).
- [22] O. Viehmann, J. von Delft and F. Marquardt, *Observing the Nonequilibrium Dynamics of the Quantum Transverse-Field Ising Chain in Circuit QED*, Phys. Rev. Lett. **110**, 030601 (2013).
- [23] L.-H. Du, J. Q. You and L. Tian, *Superconducting circuit probe for analog quantum simulators*, Phys. Rev. A **92**, 012330 (2015).
- [24] J.-M. Reiner, M. Marthaler, J. Braumüller *et al.*, *Emulating the one-dimensional Fermi-Hubbard model by a double chain of qubits*, Phys. Rev. A **94**, 032338 (2016).
- [25] M. Greiner, O. Mandel, T. Esslinger *et al.*, *Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms*, Nature **415**, 39–44 (2002).
- [26] A. Friedenauer, H. Schmitz, J. T. Glueckert *et al.*, *Simulating a quantum magnet with trapped ions*, Nat. Phys. **4**, 757 (2008).
- [27] R. Gerritsma, G. Kirchmair, F. Zähringer *et al.*, *Quantum simulation of the Dirac equation*, Nature **463**, 68 (2010).
- [28] K. Kim, M.-S. Chang, S. Korenblit *et al.*, *Quantum simulation of frustrated Ising spins with trapped ions*, Nature **465**, 590 (2010).
- [29] R. Islam, E. E. Edwards, K. Kim *et al.*, *Onset of a quantum phase transition with a trapped ion quantum simulator*, Nat. Commun. **2**, 377 (2011).
- [30] H. S. Nguyen, D. Gerace, I. Carusotto *et al.*, *Acoustic Black Hole in a Stationary Hydrodynamic Flow of Microcavity Polaritons*, Phys. Rev. Lett. **114**, 036402 (2015).
- [31] J. Steinhauer, *Observation of quantum Hawking radiation and its entanglement in an analogue black hole*, Nat. Phys. **12**, 959 (2016).
- [32] H. Labuhn, D. Barredo, S. Ravets *et al.*, *Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models*, Nature **534**, 667 (2016).
- [33] H. Bernien, S. Schwartz, A. Keesling *et al.*, *Probing many-body dynamics on a 51-atom quantum simulator*, Nature **551**, 579 (2017).
- [34] P. Roushan, C. Neill, J. Tangpanitanon *et al.*, *Spectroscopic signatures of localization with interacting photons in superconducting qubits*, Science **358**, 1175 (2017).
- [35] J. Zhang, G. Pagano, P. W. Hess *et al.*, *Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator*, Nature **551**, 601 (2017).
- [36] X. Zhang, K. Zhang, Y. Shen *et al.*, *Experimental quantum simulation of fermion-antifermion scattering via boson exchange in a trapped ion*, Nat. Commun. **9**, 195 (2018).
- [37] S. Lloyd, *Universal Quantum Simulators*, Science **273**, 1073 (1996).

- [38] G. Ortiz, J. E. Gubernatis, E. Knill *et al.*, *Quantum algorithms for fermionic simulations*, Phys. Rev. A **64**, 022319 (2001).
- [39] G. Benenti, G. Casati, S. Montangero *et al.*, *Efficient Quantum Computing of Complex Dynamics*, Phys. Rev. Lett. **87**, 227901 (2001).
- [40] R. Somma, G. Ortiz, J. E. Gubernatis *et al.*, *Simulating physical phenomena by quantum networks*, Phys. Rev. A **65**, 042323 (2002).
- [41] G. Benenti, G. Casati, S. Montangero *et al.*, *Dynamical localization simulated on a few-qubit quantum computer*, Phys. Rev. A **67**, 052312 (2003).
- [42] G. Benenti, G. Casati and S. Montangero, *Quantum computing and information extraction for a dynamical quantum system*, Quantum Inf. Process. **3**, 273 (2004).
- [43] S. Montangero, *Dynamically localized systems: Exponential sensitivity of entanglement and efficient quantum simulations*, Phys. Rev. A **70**, 032311 (2004).
- [44] M. K. Henry, J. Emerson, R. Martinez *et al.*, *Localization in the quantum sawtooth map emulated on a quantum-information processor*, Phys. Rev. A **74**, 062317 (2006).
- [45] F. Verstraete, J. I. Cirac and J. I. Latorre, *Quantum circuits for strongly correlated quantum systems*, Phys. Rev. A **79**, 032316 (2009).
- [46] H. Weimer, M. Müller, I. Lesanovsky *et al.*, *A Rydberg quantum simulator*, Nat. Phys. **6**, 382 (2010).
- [47] P. Santini, S. Carretta, F. Troiani *et al.*, *Molecular Nanomagnets as Quantum Simulators*, Phys. Rev. Lett. **107**, 230502 (2011).
- [48] J. Casanova, A. Mezzacapo, L. Lamata *et al.*, *Quantum Simulation of Interacting Fermion Lattice Models in Trapped Ions*, Phys. Rev. Lett. **108**, 190502 (2012).
- [49] A. Mezzacapo, J. Casanova, L. Lamata *et al.*, *Digital Quantum Simulation of the Holstein Model in Trapped Ions*, Phys. Rev. Lett. **109**, 200501 (2012).
- [50] S. P. Jordan, K. S. M. Lee and J. Preskill, *Quantum Algorithms for Quantum Field Theories*, Science **336**, 1130 (2012).
- [51] S. Raesi, N. Wiebe and B. C. Sanders, *Quantum-circuit design for efficient simulations of many-body quantum dynamics*, New J. Phys. **14**, 103017 (2012).
- [52] P. Hauke, D. Marcos, M. Dalmonte *et al.*, *Quantum Simulation of a Lattice Schwinger Model in a Chain of Trapped Ions*, Phys. Rev. X **3**, 041018 (2013).
- [53] U. Las Heras, A. Mezzacapo, L. Lamata *et al.*, *Digital Quantum Simulation of Spin Systems in Superconducting Circuits*, Phys. Rev. Lett. **112**, 200501 (2014).
- [54] A. Chiesa, P. Santini, D. Gerace *et al.*, *Digital quantum simulators in a scalable architecture of hybrid spin-photon qubits*, Sci. Rep. **5**, 16036 (2015).
- [55] L. García-Álvarez, I. L. Egusquiza, L. Lamata *et al.*, *Digital Quantum Simulation of Minimal AdS / CFT*, Phys. Rev. Lett. **119**, 040501 (2017).
- [56] Z. Jiang, K. Sung, K. Kechedzhi *et al.*, *Quantum Algorithms to Simulate Many-Body Physics of Correlated Fermions*, Phys. Rev. Appl. **9**, 044036 (2018).
- [57] I. D. Kivlichan, J. McClean, N. Wiebe *et al.*, *Quantum Simulation of Electronic Structure with Linear Depth and Connectivity*, Phys. Rev. Lett. **120**, 110501 (2018).

- [58] D. P. DiVincenzo, *The Physical Implementation of Quantum Computation*, Fortschr. Phys. **48**, 771 (2000).
- [59] M. Troyer and U.-J. Wiese, *Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations*, Phys. Rev. Lett. **94**, 170201 (2005).
- [60] R. Barends, J. Kelly, A. Megrant *et al.*, *Superconducting quantum circuits at the surface code threshold for fault tolerance*, Nature **508**, 500 (2014).
- [61] E. A. Martinez, C. A. Muschik, P. Schindler *et al.*, *Real-time dynamics of lattice gauge theories with a few-qubit quantum computer*, Nature **534**, 516 (2016).
- [62] N. Klco, E. F. Dumitrescu, A. J. McCaskey *et al.*, *Quantum-classical computation of Schwinger model dynamics using quantum computers*, Phys. Rev. A **98**, 032331 (2018).
- [63] A. Roggero and J. Carlson, *Linear Response on a Quantum Computer*, arXiv:1804.01505 (2018).
- [64] A. Mezzacapo, U. Las Heras, J. S. Pedernales *et al.*, *Digital Quantum Rabi and Dicke Models in Superconducting Circuits*, Sci. Rep. **4**, 7482 (2015).
- [65] I. Buluta and F. Nori, *Quantum Simulators*, Science **326**, 108–111 (2009).
- [66] I. M. Georgescu, S. Ashhab and F. Nori, *Quantum simulation*, Rev. Mod. Phys. **86**, 153 (2014).
- [67] B. C. Sanders, *Efficient algorithms for universal quantum simulation*, in G. W. Dueck and D. M. Miller (eds.), *Reversible Computation*, p. 1, Springer Berlin Heidelberg, Berlin, Heidelberg (2013).
- [68] R. Blatt and C. F. Roos, *Quantum simulations with trapped ions*, Nat. Phys. **8**, 277 (2012).
- [69] I. Bloch, J. Dalibard and S. Nascimbène, *Quantum simulations with ultracold quantum gases*, Nat. Phys. **8**, 267–276 (2012).
- [70] A. Aspuru-Guzik and P. Walther, *Photonic quantum simulators*, Nat. Phys. **8**, 285–291 (2012).
- [71] A. A. Houck, H. E. Türeci and J. Koch, *On-chip quantum simulation with superconducting circuits*, Nat. Phys. **8**, 292 (2012).
- [72] G. Wendin, *Quantum information processing with superconducting circuits: a review*, Rep. Prog. Phys. **80**, 106001 (2017).
- [73] L. Lamata, A. Parra-Rodriguez, M. Sanz *et al.*, *Digital-analog quantum simulations with superconducting circuits*, Adv. Phys.: X **3**, 1457981 (2018).
- [74] C. Monroe and J. Kim, *Scaling the Ion Trap Quantum Processor*, Science **339**, 1164 (2013).
- [75] P. Schindler, D. Nigg, T. Monz *et al.*, *A quantum information processor with trapped ions*, New J. Phys. **15**, 123012 (2013).
- [76] J. Clarke and F. K. Wilhelm, *Superconducting quantum bits*, Nature **453**, 1031–1042 (2008).
- [77] R. J. Schoelkopf and S. M. Girvin, *Wiring up quantum systems*, Nature **451**, 664–669 (2008).
- [78] M. H. Devoret and R. J. Schoelkopf, *Superconducting Circuits for Quantum Information: An Outlook*, Science **339**, 1169 (2013).
- [79] A. Acín, I. Bloch, H. Buhrman *et al.*, *The quantum technologies roadmap: a European community view*, New J. Phys. **20**, 080201 (2018).

- [80] E. Pednault, J. A. Gunnels, G. Nannicini *et al.*, *Breaking the 49-Qubit Barrier in the Simulation of Quantum Circuits*, arXiv:1710.05867 (2017).
- [81] S. Boixo, S. V. Isakov, V. N. Smelyanskiy *et al.*, *Characterizing quantum supremacy in near-term devices*, Nat. Phys. **14**, 595–600 (2018).
- [82] B. Villalonga, D. Lyakh, S. Boixo *et al.*, *Establishing the Quantum Supremacy Frontier with a 281 Pflop/s Simulation*, arXiv:1905.00444 (2019).
- [83] F. Arute, K. Arya, R. Babbush *et al.*, *Quantum supremacy using a programmable superconducting processor*, Nature **574**, 505–510 (2019).
- [84] P. Schindler, J. T. Barreiro, T. Monz *et al.*, *Experimental Repetitive Quantum Error Correction*, Science **332**, 1059 (2011).
- [85] H. You, M. R. Geller and P. C. Stancil, *Simulating the transverse Ising model on a quantum computer: Error correction with the surface code*, Phys. Rev. A **87**, 032341 (2013).
- [86] A. Córcoles, E. Magesan, S. J. Srinivasan *et al.*, *Demonstration of a quantum error detection code using a square lattice of four superconducting qubits*, Nat. Commun. **6**, 6979 (2015).
- [87] S. Woerner and D. J. Egger, *Quantum risk analysis*, npj Quantum Inf. **5**, 15 (2019).
- [88] A. Martin, B. Candelas, Á. Rodríguez-Rozas *et al.*, *Towards Pricing Financial Derivatives with an IBM Quantum Computer*, arXiv:1904.05803 (2019).
- [89] J. Biamonte, P. Wittek, N. Pancotti *et al.*, *Quantum machine learning*, Nature **549**, 195 (2017).
- [90] C. Moler and C. Van Loan, *Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later*, SIAM Review **45**, 3–49 (2003).
- [91] A. Barenco, C. H. Bennett, R. Cleve *et al.*, *Elementary gates for quantum computation*, Phys. Rev. A **52**, 3457 (1995).
- [92] F. Tacchino, A. Chiesa, M. D. LaHaye *et al.*, *Electromechanical quantum simulators*, Phys. Rev. B **97**, 214302 (2018).
- [93] R. Barends, L. Lamata, J. Kelly *et al.*, *Digital quantum simulation of fermionic models with a superconducting circuit*, Nat. Commun. **6**, 7654 (2015).
- [94] A. Kandala, A. Mezzacapo, K. Temme *et al.*, *Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets*, Nature **549**, 242 (2017).
- [95] E. F. Dumitrescu, A. J. McCaskey, G. Hagen *et al.*, *Cloud Quantum Computing of an Atomic Nucleus*, Phys. Rev. Lett. **120**, 210501 (2018).
- [96] P. Jordan and E. Wigner, *Über das Paulische äquivalenzverbot*, Zeitschrift für Physik **47**, 631–651 (1928).
- [97] R. A. Bari, *Classical Linear-Chain Hubbard Model: Metal-Insulator Transition*, Phys. Rev. B **7**, 4318 (1973).
- [98] N. M. Linke, D. Maslov, M. Roetteler *et al.*, *Experimental comparison of two quantum computing architectures*, PNAS **114**, 3305 (2017).
- [99] A. W. Cross, L. S. Bishop, S. Sheldon *et al.*, *Validating quantum computers using randomized model circuits*, arXiv:1811.12926 (2018).
- [100] D. C. McKay, C. J. Wood, S. Sheldon *et al.*, *Efficient Z gates for quantum computing*, Phys. Rev. A **96**, 022330 (2017).

- [101] J. M. Chow, J. M. Gambetta, A. D. Córcoles *et al.*, *Universal Quantum Gate Set Approaching Fault-Tolerant Thresholds with Superconducting Qubits*, Phys. Rev. Lett. **109**, 060501 (2012).
- [102] G. S. Paraoanu, *Microwave-induced coupling of superconducting qubits*, Phys. Rev. B **74**, 140504 (2006).
- [103] C. Rigetti and M. Devoret, *Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies*, Phys. Rev. B **81**, 134507 (2010).
- [104] J. M. Chow, A. D. Córcoles, J. M. Gambetta *et al.*, *Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits*, Phys. Rev. Lett. **107**, 080502 (2011).
- [105] S. Sheldon, E. Magesan, J. M. Chow *et al.*, *Procedure for systematically tuning up cross-talk in the cross-resonance gate*, Phys. Rev. A **93**, 060302 (2016).
- [106] Y. Salathé, M. Mondal, M. Oppliger *et al.*, *Digital Quantum Simulation of Spin Models with Circuit Quantum Electrodynamics*, Phys. Rev. X **5**, 021027 (2015).
- [107] D. C. McKay, S. Filipp, A. Mezzacapo *et al.*, *Universal Gate for Fixed-Frequency Qubits via a Tunable Bus*, Phys. Rev. Appl. **6**, 064007 (2016).
- [108] N. Schuch and J. Siewert, *Natural two-qubit gate for quantum computation using the XY interaction*, Phys. Rev. A **67**, 032301 (2003).
- [109] P. Krantz, M. Kjaergaard, F. Yan *et al.*, *A Quantum Engineer's Guide to Superconducting Qubits*, arXiv:1904.06560 (2019).
- [110] S. Carretta, A. Chiesa, F. Troiani *et al.*, *Quantum Information Processing with Hybrid Spin-Photon Qubit Encoding*, Phys. Rev. Lett. **111**, 110501 (2013).
- [111] A. Chiesa, D. Gerace, F. Troiani *et al.*, *Robustness of quantum gates with hybrid spin-photon qubits in superconducting resonators*, Phys. Rev. A **89**, 052308 (2014).
- [112] M. Reagor, C. B. Osborn, N. Tezak *et al.*, *Demonstration of universal parametric entangling gates on a multi-qubit lattice*, Sci. Adv. **4**, eaao3603 (2018).
- [113] J. A. Jones, *Robust Ising gates for practical quantum computation*, Phys. Rev. A **67**, 012317 (2003).
- [114] P. Schindler, M. Müller, D. Nigg *et al.*, *Quantum simulation of dynamical maps with trapped ions*, Nat. Phys. **9**, 361 (2013).
- [115] B. P. Lanyon, C. Hempel, D. Nigg *et al.*, *Universal Digital Quantum Simulation with Trapped Ions*, Science **334**, 57 (2011).
- [116] K. Mølmer and A. Sørensen, *Multiparticle Entanglement of Hot Trapped Ions*, Phys. Rev. Lett. **82**, 1835 (1999).
- [117] A. E. Webb, S. C. Webster, S. Collingbourne *et al.*, *Resilient Entangling Gates for Trapped Ions*, Phys. Rev. Lett. **121**, 180501 (2018).
- [118] Y. Shapira, R. Shaniv, T. Manovitz *et al.*, *Robust Entanglement Gates for Trapped-Ion Qubits*, Phys. Rev. Lett. **121**, 180502 (2018).
- [119] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov *et al.*, *Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions*, Phys. Rev. A **98**, 022322 (2018).
- [120] M. Müller, K. Hammerer, Y. L. Zhou *et al.*, *Simulating open quantum systems: from many-body interactions to stabilizer pumping*, New J. Phys. **13**, 085007 (2011).

- [121] N. Hatano and M. Suzuki, *Finding Exponential Product Formulas of Higher Orders*, in A. Das and B. K. Chakrabarti (eds.), *Quantum Annealing and Other Optimization Methods*, pp. 37–68, Springer Berlin Heidelberg, Berlin, Heidelberg (2005).
- [122] R. Kubo, *Statistical-Mechanical Theory of Irreversible Processes. I. General Theory and Simple Applications to Magnetic and Conduction Problems*, Journal of the Physical Society of Japan **12**, 570–586 (1957).
- [123] A. L. Fetter and J. D. Walecka, *Quantum theory of many-particle systems*, International series in pure and applied physics, McGraw-Hill, S. Francisco (1971).
- [124] A. Chiesa, F. Tacchino, M. Grossi *et al.*, *Quantum hardware simulating four-dimensional inelastic neutron scattering*, Nat. Phys. **15**, 455 (2019).
- [125] J. S. Pedernales, R. Di Candia, I. L. Egusquiza *et al.*, *Efficient Quantum Algorithm for Computing n -time Correlation Functions*, Phys. Rev. Lett. **113**, 020505 (2014).
- [126] Cleve R., Ekert A., Macchiavello C. *et al.*, *Quantum algorithms revisited*, Proc. Royal Soc. A **454**, 339–354 (1998).
- [127] G. Vidal and C. M. Dawson, *Universal quantum circuit for two-qubit transformations with three controlled-NOT gates*, Phys. Rev. A **69**, 010301 (2004).
- [128] J. Ferrando-Soria, S. A. Magee, A. Chiesa *et al.*, *Switchable Interaction in Molecular Double Qubits*, Chem **1**, 727–752 (2016).
- [129] J.-M. Raimond and S. Haroche, *Exploring the quantum: atoms, cavities and photons*, Oxford University Press, Oxford (2006).
- [130] D. J. Wineland, *Nobel Lecture: Superposition, entanglement, and raising Schrödinger’s cat*, Rev. Mod. Phys. **85**, 1103–1114 (2013).
- [131] W. Paul, *Electromagnetic traps for charged and neutral particles*, Rev. Mod. Phys. **62**, 531–540 (1990).
- [132] R. Blatt and D. Wineland, *Entangled states of trapped atomic ions*, Nature **453**, 1008–1015 (2008).
- [133] C. D. Bruzewicz, J. Chiaverini, R. McConnell *et al.*, *Trapped-Ion Quantum Computing: Progress and Challenges*, arXiv:1904.04178 (2019).
- [134] K. Wright, K. M. Beck, S. Debnath *et al.*, *Benchmarking an 11-qubit quantum computer*, arxiv:1903.08181 (2019).
- [135] D. Leibfried, R. Blatt, C. Monroe *et al.*, *Quantum dynamics of single trapped ions*, Rev. Mod. Phys. **75**, 281–324 (2003).
- [136] A. Bautista-Salvador, G. Zarantonello, H. Hahn *et al.*, *Multilayer ion trap technology for scalable quantum computing and quantum simulation*, New J. Phys. **21**, 043011 (2019).
- [137] C. D. Bruzewicz, R. McConnell, J. Chiaverini *et al.*, *Scalable loading of a two-dimensional trapped-ion array*, Nature Communications **7**, 13005 (2016).
- [138] N. Friis, O. Marty, C. Maier *et al.*, *Observation of Entangled States of a Fully Controlled 20-Qubit System*, Phys. Rev. X **8**, 021012 (2018).
- [139] C. Hempel, C. Maier, J. Romero *et al.*, *Quantum chemistry calculations on a trapped-ion quantum simulator*, Phys. Rev. X **8**, 031022 (2018).
- [140] C. J. Ballance, T. P. Harty, N. M. Linke *et al.*, *High-fidelity quantum logic gates using trapped-ion hyperfine qubits*, Phys. Rev. Lett. **117**, 060504 (2016).

- [141] S. Debnath, N. M. Linke, C. Figgatt *et al.*, *Demonstration of a small programmable quantum computer with atomic qubits*, *Nature* **536**, 63 (2016).
- [142] J. I. Cirac and P. Zoller, *Quantum computations with cold trapped ions*, *Phys. Rev. Lett.* **74**, 4091–4094 (1995).
- [143] T. P. Harty, D. T. C. Allcock, C. J. Ballance *et al.*, *High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit*, *Phys. Rev. Lett.* **113**, 220501 (2014).
- [144] J. P. Gaebler, T. R. Tan, Y. Lin *et al.*, *High-fidelity universal gate set for ${}^9\text{Be}^+$ ion qubits*, *Phys. Rev. Lett.* **117**, 060505 (2016).
- [145] V. M. Schäfer, C. J. Ballance, K. Thirumalai *et al.*, *Fast quantum logic gates with trapped-ion qubits*, *Nature* **555**, 75–78 (2018).
- [146] A. H. Myerson, D. J. Szwer, S. C. Webster *et al.*, *High-fidelity readout of trapped-ion qubits*, *Phys. Rev. Lett.* **100**, 200502 (2008).
- [147] A. Bermudez, X. Xu, R. Nigmatullin *et al.*, *Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation*, *Phys. Rev. X* **7**, 041061 (2017).
- [148] Y. Wang, M. Um, J. Zhang *et al.*, *Single-qubit quantum memory exceeding ten-minute coherence time*, *Nature Photonics* **11**, 646–650 (2017).
- [149] A. Erhard, J. J. Wallman, L. Postler *et al.*, *Characterizing large-scale quantum computers via cycle benchmarking*, arxiv:1902.08543 (2019).
- [150] C. Kokail, C. Maier, R. van Bijnen *et al.*, *Self-verifying variational quantum simulation of lattice models*, *Nature* **569**, 355 (2019).
- [151] J. R. McClean, J. Romero, R. Babbush *et al.*, *The theory of variational hybrid quantum-classical algorithms*, *New J. Phys.* **18**, 023023 (2016).
- [152] Y. Nam, J.-S. Chen, N. C. Pienti *et al.*, *Ground-state energy estimation of the water molecule on a trapped ion quantum computer*, arxiv:1902.10171 (2019).
- [153] O. Shehab, K. A. Landsman, Y. Nam *et al.*, *Toward convergence of effective field theory simulations on digital quantum computers*, arXiv:1904.04338 (2019).
- [154] A. Blais, R.-S. Huang, A. Wallraff *et al.*, *Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation*, *Phys. Rev. A* **69**, 062320 (2004).
- [155] A. Wallraff, D. I. Schuster, A. Blais *et al.*, *Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics*, *Nature* **431**, 162–167 (2004).
- [156] J. Koch, T. M. Yu, J. M. Gambetta *et al.*, *Charge-insensitive qubit design derived from the Cooper pair box*, *Phys. Rev. A* **76**, 042319 (2007).
- [157] J. Gambetta, W. A. Braff, A. Wallraff *et al.*, *Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement*, *Phys. Rev. A* **76**, 012325 (2007).
- [158] J. Majer, J. M. Chow, J. M. Gambetta *et al.*, *Coupling superconducting qubits via a cavity bus*, *Nature* **449**, 443–447 (2007).
- [159] J. Gambetta, A. Blais, M. Boissonneault *et al.*, *Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect*, *Phys. Rev. A* **77**, 012112 (2008).
- [160] M. Mariantoni, H. Wang, T. Yamamoto *et al.*, *Implementing the Quantum von Neumann Architecture with Superconducting Circuits*, *Science* **334**, 61 (2011).

- [161] C. Rigetti, J. M. Gambetta, S. Poletto *et al.*, *Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms*, Phys. Rev. B **86**, 100506(R) (2012).
- [162] A. Nersisyan, S. Poletto, N. Alidoust *et al.*, *Manufacturing low dissipation superconducting quantum processors*, arXiv:1901.08042 (2019).
- [163] J. Kelly, R. Barends, A. G. Fowler *et al.*, *State preservation by repetitive error detection in a superconducting quantum circuit*, Nature **519**, 66–69 (2015).
- [164] R. Barends, A. Shabani, L. Lamata *et al.*, *Digitized adiabatic quantum computing with a superconducting circuit*, Nature **534**, 222 (2016).
- [165] L. DiCarlo, J. M. Chow, J. Gambetta *et al.*, *Demonstration of two-qubit algorithms with a superconducting quantum processor*, Nature **460**, 240–244 (2009).
- [166] J. M. Gambetta, J. M. Chow and M. Steffen, *Building logical qubits in a superconducting quantum computing system*, npj Quantum Inf. **3**, 2 (2017).
- [167] M. Ganzhorn, D. Egger, P. Barkoutsos *et al.*, *Gate-efficient simulation of molecular eigenstates on a quantum computer*, Phys. Rev. A **11**, 044092 (2019).
- [168] U. Las Heras, L. García-Álvarez, A. Mezzacapo *et al.*, *Fermionic models with superconducting circuits*, EPJ Quantum Tech. **2**, 8 (2015).
- [169] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan *et al.*, *Scalable Quantum Simulation of Molecular Energies*, Phys. Rev. X **6**, 031007 (2016).
- [170] K. Temme, S. Bravyi and J. M. Gambetta, *Error Mitigation for Short-Depth Quantum Circuits*, Phys. Rev. Lett. **119**, 180509 (2017).
- [171] S. McArdle, X. Yuan and S. Benjamin, *Error-mitigated digital quantum simulation*, Phys. Rev. Lett. **122**, 180501 (2019).
- [172] A. Kandala, K. Temme, A. D. Córcoles *et al.*, *Error mitigation extends the computational reach of a noisy quantum processor*, Nature **567**, 491 (2019).
- [173] T. Roy, S. Hazra, S. Kundu *et al.*, *A programmable three-qubit superconducting processor with all-to-all connectivity*, arXiv:1809.00668 (2018).
- [174] P. V. Klimov, J. Kelly, Z. Chen *et al.*, *Fluctuations of energy-relaxation times in superconducting qubits*, Phys. Rev. Lett. **121**, 090502 (2018).
- [175] N. Moll, P. Barkoutsos, L. S. Bishop *et al.*, *Quantum optimization using variational algorithms on near-term quantum devices*, Quantum Sci. Tech. **3**, 030503 (2018).
- [176] K. X. Wei, I. Lauer, S. Srinivasan *et al.*, *Verifying Multipartite Entangled GHZ States via Multiple Quantum Coherences*, arXiv:1905.05720 (2019).
- [177] J. M. Gambetta, *Benchmarking NISQ-era quantum processors, Invited Talk Y34.00004 at APS March Meeting* (2019).
- [178] C. Figgatt, A. Ostrander, N. M. Linke *et al.*, *Parallel entangling operations on a universal ion-trap quantum computer*, Nature **572**, 368–372 (2019).
- [179] Y. Lu, S. Zhang, K. Zhang *et al.*, *Global entangling gates on arbitrary ion qubits*, Nature **572**, 363–367 (2019).
- [180] Y. Cao, J. Romero, J. P. Olson *et al.*, *Quantum Chemistry in the Age of Quantum Computing*, arXiv:1812.09976 (2018).

- [181] F. Fillion-Gourdeau, S. MacLean and R. Laflamme, *Algorithm for the solution of the Dirac equation on digital quantum computers*, Phys. Rev. A **95**, 042343 (2017).
- [182] T. V. Zache, F. Hebenstreit, F. Jendrzejewski *et al.*, *Quantum simulation of lattice gauge theories using Wilson fermions*, Quantum Sci. Technol. **3**, 034010 (2018).
- [183] C. W. Bauer, W. A. de Jong, B. Nachman *et al.*, *A quantum algorithm for high energy physics simulations*, arXiv:1904.03196 (2019).
- [184] M. Kliesch, T. Barthel, C. Gogolin *et al.*, *Dissipative Quantum Church-Turing Theorem*, Phys. Rev. Lett. **107**, 120501 (2011).
- [185] H. Wang, S. Ashhab and F. Nori, *Quantum algorithm for simulating the dynamics of an open quantum system*, Phys. Rev. A **83**, 062317 (2011).
- [186] R. Di Candia, J. S. Pedernales, A. del Campo *et al.*, *Quantum Simulation of Dissipative Processes without Reservoir Engineering*, Sci. Rep. **5**, 09981 (2015).
- [187] A. Chenu, M. Beau, J. Cao *et al.*, *Quantum Simulation of Generic Many-Body Open System Dynamics Using Classical Noise*, Phys. Rev. Lett. **118**, 140403 (2017).
- [188] R. Cleve and C. Wang, *Efficient Quantum Algorithms for Simulating Lindblad Evolution*, in *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, vol. 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, p. 17:1, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2017).
- [189] H. Lamm and S. Lawrence, *Simulation of Nonequilibrium Dynamics on a Quantum Computer*, Phys. Rev. Lett. **121**, 170501 (2018).
- [190] A. Smith, M. S. Kim, F. Pollmann *et al.*, *Simulating quantum many-body dynamics on a current digital quantum computer*, arXiv:1906.06343 (2019).
- [191] M. J. S. Beach, R. G. Melko, T. Grover *et al.*, *Making Trotters Sprint: A Variational Imaginary Time Ansatz for Quantum Many-body Systems*, arXiv:1904.00019 (2019).
- [192] M. Motta, C. Sun, A. T. K. Tan *et al.*, *Quantum Imaginary Time Evolution, Quantum Lanczos, and Quantum Thermal Averaging*, arXiv:1901.07653 (2019).
- [193] A. M. Childs and Y. Su, *Nearly Optimal Lattice Simulation by Product Formulas*, Phys. Rev. Lett. **123**, 050503 (2019).
- [194] U. Las Heras, U. Alvarez-Rodriguez, E. Solano *et al.*, *Genetic Algorithms for Digital Quantum Simulations*, Phys. Rev. Lett. **116**, 230504 (2016).
- [195] Y. Li and S. C. Benjamin, *Efficient Variational Quantum Simulator Incorporating Active Error Minimization*, Phys. Rev. X **7**, 021050 (2017).
- [196] G. Aleksandrowicz, T. Alexander, P. Barkoutsos *et al.*, *Qiskit: An open-source framework for quantum computing* (2019).
- [197] A. W. Cross, L. S. Bishop, J. A. Smolin *et al.*, *Open Quantum Assembly Language*, arXiv:1707.03429 (2017).
- [198] F. Motzoi, J. M. Gambetta, P. Rebentrost *et al.*, *Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits*, Phys. Rev. Lett. **103**, 110501 (2009).
- [199] J. M. Chow, L. DiCarlo, J. M. Gambetta *et al.*, *Optimized driving of superconducting artificial atoms for improved single-qubit gates*, Phys. Rev. A **82**, 040305 (2010).
- [200] J. M. Chow, J. M. Gambetta, A. W. Cross *et al.*, *Microwave-activated conditional-phase gate for superconducting qubits*, New J. Phys. **15**, 115012 (2013).

- [201] M. Fingerhuth, T. Babej and P. Wittek, *Open source software in quantum computing*, PLOS ONE **13**, e0208561 (2018).
- [202] F. Tacchino, A. Chiesa, M. D. LaHaye *et al.*, *Digital Quantum Simulations of Spin Models on Hybrid Platform and Near-Term Quantum Processors*, Proceedings **12**, 24 (2019).
- [203] D. Gatteschi, R. Sessoli and J. Villain, *Molecular Nanomagnets*, Oxford University Press, New York, USA (2006).
- [204] J. Johansson, P. Nation and F. Nori, *QuTiP: An open-source Python framework for the dynamics of open quantum systems*, Comput. Phys. Commun. **183**, 1760–1772 (2012).
- [205] J. Johansson, P. Nation and F. Nori, *QuTiP 2: A Python framework for the dynamics of open quantum systems*, Comput. Phys. Commun. **184**, 1234–1240 (2013).
- [206] A. A. Clerk, M. H. Devoret, S. M. Girvin *et al.*, *Introduction to quantum noise, measurement, and amplification*, Rev. Mod. Phys. **82**, 1155–1208 (2010).
- [207] D. Sank, E. Jeffrey, J. Mutus *et al.*, *Fast scalable state measurement with superconducting qubits*, Phys. Rev. Lett. **112**, 190504 (2014).
- [208] A. Furrer and O. Waldmann, *Magnetic cluster excitations*, Rev. Mod. Phys. **85**, 367–420 (2013).
- [209] C. Schröder, H. Nojiri, J. Schnack *et al.*, *Competing Spin Phases in Geometrically Frustrated Magnetic Molecules*, Phys. Rev. Lett. **94**, 017205 (2005).
- [210] A. M. Todea, A. Merca, H. Bögge *et al.*, *Extending the $\{(Mo)Mo_5\}12m30$ Capsule Keplerate Sequence: A $\{Cr_{30}\}$ Cluster of $S=3/2$ Metal Centers with a $\{Na(H_2O)_{12}\}$ Encapsulate*, Angew. Chem., Int. Ed. **46**, 6106–6110 (2007).
- [211] G. F. S. Whitehead, F. Moro, G. A. Timco *et al.*, *A Ring of Rings and Other Multicomponent Assemblies of Cages*, Angewandte Chemie International Edition **52**, 9932–9935 (2013).
- [212] A. Baniodeh, N. Magnani, Y. Lan *et al.*, *High spin cycles: topping the spin record for a single molecule verging on quantum criticality*, npj Quantum Mater. **3**, 10 (2018).
- [213] M. L. Baker, T. Guidi, S. Carretta *et al.*, *Spin dynamics of molecular nanomagnets unravelled at atomic scale by four-dimensional inelastic neutron scattering*, Nat. Phys. **8**, 906–911 (2012).
- [214] E. Garlatti, T. Guidi, S. Ansbro *et al.*, *Portraying entanglement between molecular qubits with four-dimensional inelastic neutron scattering*, Nat. Commun. **8**, 14543 (2017).
- [215] A. Chiesa, T. Guidi, S. Carretta *et al.*, *Magnetic Exchange Interactions in the Molecular Nanomagnet Mn 12*, Phys. Rev. Lett. **119**, 217202 (2017).
- [216] S. W. Lovesey, *Theory of Neutron Scattering from Condensed Matter*, Oxford University Press, Oxford, UK (2003).
- [217] S. S. Eaton, G. R. Eaton and C. K. Chang, *Synthesis and geometry determination of cofacial diporphyrins. EPR spectroscopy of dicopper diporphyrins in frozen solution*, Journal of the American Chemical Society **107**, 3177–3184 (1985).
- [218] S. Hill and W. K. Wootters, *Entanglement of a Pair of Quantum Bits*, Phys. Rev. Lett. **78**, 5022–5025 (1997).
- [219] W. K. Wootters, *Entanglement of Formation of an Arbitrary State of Two Qubits*, Phys. Rev. Lett. **80**, 2245–2248 (1998).
- [220] W. K. Wootters, *Entanglement of formation and concurrence*, Quantum Inf. Comput. **1**, 27–44 (2001).

- [221] D. Loss and D. P. DiVincenzo, *Quantum computation with quantum dots*, Phys. Rev. A **57**, 120 (1998).
- [222] L. M. K. Vandersypen and M. A. Eriksson, *Quantum computing with semiconductor spins*, Physics Today **72**, 38–45 (2019).
- [223] J. M. Elzerman, R. Hanson, L. H. Willems van Beveren *et al.*, *Single-shot read-out of an individual electron spin in a quantum dot*, Nature **430**, 431–435 (2004).
- [224] J. R. Petta, A. C. Johnson, J. M. Taylor *et al.*, *Coherent Manipulation of Coupled Electron Spins in Semiconductor Quantum Dots*, Science **309**, 2180 (2005).
- [225] R. Hanson, L. P. Kouwenhoven, J. R. Petta *et al.*, *Spins in few-electron quantum dots*, Rev. Mod. Phys. **79**, 1217–1265 (2007).
- [226] D. M. Zajac, A. J. Sigillito, M. Russ *et al.*, *Resonantly driven CNOT gate for electron spins*, Science **359**, 439 (2018).
- [227] M. Veldhorst, C. H. Yang, J. C. C. Hwang *et al.*, *A two-qubit logic gate in silicon*, Nature **526**, 410–414 (2015).
- [228] T. F. Watson, S. G. J. Philips, E. Kawakami *et al.*, *A programmable two-qubit quantum processor in silicon*, Nature **555**, 633 (2018).
- [229] A. Noiri, T. Nakajima, J. Yoneda *et al.*, *A fast quantum interface between different spin qubit encodings*, Nat. Commun. **9**, 5066 (2018).
- [230] W. Huang, C. H. Yang, K. W. Chan *et al.*, *Fidelity benchmarks for two-qubit gates in silicon*, Nature (2019).
- [231] D. D. Awschalom, L. C. Bassett, A. S. Dzurak *et al.*, *Quantum Spintronics: Engineering and Manipulating Atom-Like Spins in Semiconductors*, Science **339**, 1174 (2013).
- [232] A. Morello, J. J. Pla, F. A. Zwanenburg *et al.*, *Single-shot readout of an electron spin in silicon*, Nature **467**, 687–691 (2010).
- [233] J. J. Pla, K. Y. Tan, J. P. Dehollain *et al.*, *A single-atom electron spin qubit in silicon*, Nature **489**, 541–545 (2012).
- [234] R. Kalra, A. Laucht, C. D. Hill *et al.*, *Robust Two-Qubit Gates for Donors in Silicon Controlled by Hyperfine Interactions*, Phys. Rev. X **4**, 021044 (2014).
- [235] G. Tosi, F. A. Mohiyaddin, V. Schmitt *et al.*, *Silicon quantum processor with robust long-distance qubit couplings*, Nat. Commun. **8**, 450 (2017).
- [236] Y. He, S. K. Gorman, D. Keith *et al.*, *A two-qubit gate between phosphorus donor electrons in silicon*, Nature **571**, 371–375 (2019).
- [237] C. E. Bradley, J. Randall, M. H. Abobeih *et al.*, *A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute*, Phys. Rev. X **9**, 031045 (2019).
- [238] S. Ferretti and D. Gerace, *Single-photon nonlinear optics with Kerr-type nanostructured materials*, Phys. Rev. B **85**, 033303 (2012).
- [239] H. Flayac, D. Gerace and V. Savona, *An all-silicon single-photon source by unconventional photon blockade*, Sci. Rep. **5**, 11223 (2015).
- [240] G. Muñoz-Matutano, A. Wood, M. Johnsson *et al.*, *Emergence of quantum correlations from interacting fibre-cavity polaritons*, Nat. Mater. **18**, 213–218 (2019).

- [241] A. Delteil, T. Fink, A. Schade *et al.*, *Towards polariton blockade of confined exciton–polaritons*, Nat. Mater. **18**, 219–222 (2019).
- [242] D. Gerace, H. E. Türeci, A. Imamoglu *et al.*, *The quantum-optical Josephson interferometer*, Nat. Phys. **5**, 281 (2009).
- [243] S. Lloyd and S. L. Braunstein, *Quantum Computation over Continuous Variables*, Phys. Rev. Lett. **82**, 4 (1999).
- [244] N. C. Menicucci, P. van Loock, M. Gu *et al.*, *Universal Quantum Computation with Continuous-Variable Cluster States*, Phys. Rev. Lett. **97**, 110501 (2006).
- [245] C. Weedbrook, S. Pirandola, R. García-Patrón *et al.*, *Gaussian quantum information*, Rev. Mod. Phys. **84**, 621–669 (2012).
- [246] N. Killoran, T. R. Bromley, J. M. Arrazola *et al.*, *Continuous-variable quantum neural networks*, arXiv:1806.06871 (2018).
- [247] J. Ferrando-Soria, E. Moreno Pineda, A. Chiesa *et al.*, *A modular design of molecular qubits to implement universal quantum gates*, Nat. Commun. **7**, 11377 (2016).
- [248] M. Atzori, A. Chiesa, E. Morra *et al.*, *A two-qubit molecular architecture for electron-mediated nuclear quantum simulation*, Chem. Sci. **9**, 6183–6192 (2018).
- [249] R. Hussain, G. Allodi, A. Chiesa *et al.*, *Coherent Manipulation of a Molecular Ln-Based Nuclear Qudit Coupled to an Electron Qubit*, J. Am. Chem. Soc. **140**, 9814–9818 (2018).
- [250] H. Levine, A. Keesling, A. Omran *et al.*, *High-fidelity control and entanglement of rydberg-atom qubits*, Phys. Rev. Lett. **121**, 123603 (2018).
- [251] M. Saffman, *Quantum computing with atomic qubits and rydberg interactions: progress and challenges*, J. Phys. B: At., Mol. Opt. Phys. **49**, 202001 (2016).
- [252] Z.-L. Xiang, S. Ashhab, J. Q. You *et al.*, *Hybrid quantum circuits: Superconducting circuits interacting with other quantum systems*, Rev. Mod. Phys. **85**, 623 (2013).
- [253] G. Kurizki, P. Bertet, Y. Kubo *et al.*, *Quantum technologies with hybrid systems*, PNAS **112**, 3866 (2015).
- [254] N. Bar-Gill, L. M. Pham, A. Jarmola *et al.*, *Solid-state electronic spin coherence time approaching one second*, Nat. Commun. **4**, 1743 (2013).
- [255] K. Saeedi, S. Simmons, J. Z. Salvail *et al.*, *Room-Temperature Quantum Bit Storage Exceeding 39 Minutes Using Ionized Donors in Silicon-28*, Science **342**, 830–833 (2013).
- [256] A. Imamoglu, *Cavity QED Based on Collective Magnetic Dipole Coupling: Spin Ensembles as Hybrid Two-Level Systems*, Phys. Rev. Lett. **102**, 083602 (2009).
- [257] J. H. Wesenberg, A. Ardavan, G. A. D. Briggs *et al.*, *Quantum Computing with an Electron Spin Ensemble*, Phys. Rev. Lett. **103**, 070502 (2009).
- [258] P. Rabl, D. DeMille, J. M. Doyle *et al.*, *Hybrid Quantum Processors: Molecular Ensembles as Quantum Memory for Solid State Circuits*, Phys. Rev. Lett. **97**, 033003 (2006).
- [259] D. Petrosyan, G. Bensky, G. Kurizki *et al.*, *Reversible state transfer between superconducting qubits and atomic ensembles*, Phys. Rev. A **79**, 040304 (2009).
- [260] J. Verdú, H. Zoubi, C. Koller *et al.*, *Strong Magnetic Coupling of an Ultracold Gas to a Superconducting Waveguide Cavity*, Phys. Rev. Lett. **103**, 043603 (2009).

- [261] Y. Ping, E. M. Gauger and S. C. Benjamin, *Measurement-based quantum computing with a spin ensemble coupled to a stripline cavity*, New J. Phys. **14**, 013030 (2012).
- [262] A. Chiesa, P. Santini, D. Gerace *et al.*, *Long-lasting hybrid quantum information processing in a cavity-protection regime*, Phys. Rev. B **93**, 094432 (2016).
- [263] P.-B. Li, Z.-L. Xiang, P. Rabl *et al.*, *Hybrid Quantum Device with Nitrogen-Vacancy Centers in Diamond Coupled to Carbon Nanotubes*, Phys. Rev. Lett. **117**, 015502 (2016).
- [264] P.-B. Li, Y.-C. Liu, S.-Y. Gao *et al.*, *Hybrid Quantum Device Based on NV Centers in Diamond Nanomechanical Resonators Plus Superconducting Waveguide Cavities*, Phys. Rev. Appl. **4**, 044003 (2015).
- [265] S. D. Hogan, J. A. Agner, F. Merkt *et al.*, *Driving Rydberg-Rydberg Transitions from a Coplanar Microwave Waveguide*, Phys. Rev. Lett. **108**, 063004 (2012).
- [266] C. Hermann-Avigliano, R. C. Teixeira, T. L. Nguyen *et al.*, *Long coherence times for Rydberg qubits on a superconducting atom chip*, Phys. Rev. A **90**, 040502 (2014).
- [267] M. Poot and H. S. J. van der Zant, *Mechanical systems in the quantum regime*, Physics Reports **511**, 273–335 (2012).
- [268] M. Aspelmeyer, T. J. Kippenberg and F. Marquardt, *Cavity optomechanics*, Rev. Mod. Phys. **86**, 1391–1452 (2014).
- [269] S. Rips and M. J. Hartmann, *Quantum Information Processing with Nanomechanical Qubits*, Phys. Rev. Lett. **110**, 120503 (2013).
- [270] D. Rugar, R. Budakian, H. J. Mamin *et al.*, *Single spin detection by magnetic resonance force microscopy*, Nature **430**, 329–332 (2004).
- [271] M. Poggio and C. L. Degen, *Force-detected nuclear magnetic resonance: recent advances and future challenges*, Nanotechnology **21**, 342001 (2010).
- [272] S. Barzanjeh, S. Guha, C. Weedbrook *et al.*, *Microwave Quantum Illumination*, Phys. Rev. Lett. **114**, 080503 (2015).
- [273] A. M. Jayich, J. C. Sankey, B. M. Zwickl *et al.*, *Dispersive optomechanics: a membrane inside a cavity*, New J. Phys. **10**, 095008 (2008).
- [274] J. D. Thompson, B. M. Zwickl, A. M. Jayich *et al.*, *Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane*, Nature **452**, 72–75 (2008).
- [275] R. W. Andrews, R. W. Peterson, T. P. Purdy *et al.*, *Bidirectional and efficient conversion between microwave and optical light*, Nat. Phys. **10**, 321–326 (2014).
- [276] I. Wilson-Rae, N. Nooshi, W. Zwerger *et al.*, *Theory of Ground State Cooling of a Mechanical Oscillator Using Dynamical Backaction*, Phys. Rev. Lett. **99**, 093901 (2007).
- [277] F. Marquardt, J. P. Chen, A. A. Clerk *et al.*, *Quantum Theory of Cavity-Assisted Sideband Cooling of Mechanical Motion*, Phys. Rev. Lett. **99**, 093902 (2007).
- [278] T. A. Palomaki, J. D. Teufel, R. W. Simmonds *et al.*, *Entangling Mechanical Motion with Microwave Fields*, Science **342**, 710–713 (2013).
- [279] J. D. Teufel, D. Li, M. S. Allman *et al.*, *Circuit cavity electromechanics in the strong-coupling regime*, Nature **471**, 204–208 (2011).
- [280] W. P. Bowen and G. J. Milburn, *Quantum Optomechanics*, CRC Press, Boca Raton, FL, USA (2016).

- [281] J. D. Teufel, T. Donner, D. Li *et al.*, *Sideband cooling of micromechanical motion to the quantum ground state*, Nature **475**, 359–363 (2011).
- [282] T. A. Palomaki, J. W. Harlow, J. D. Teufel *et al.*, *Coherent state transfer between itinerant microwave fields and a mechanical oscillator*, Nature **495**, 210–214 (2013).
- [283] M. D. LaHaye, J. Suh, P. M. Echternach *et al.*, *Nanomechanical measurements of a superconducting qubit*, Nature **459**, 960–964 (2009).
- [284] A. D. O’Connell, M. Hofheinz, M. Ansmann *et al.*, *Quantum ground state and single-phonon control of a mechanical resonator*, Nature **464**, 697–703 (2010).
- [285] J.-M. Pirkkalainen, S. U. Cho, J. Li *et al.*, *Hybrid circuit cavity quantum electrodynamics with a micromechanical resonator*, Nature **494**, 211–215 (2013).
- [286] J.-M. Pirkkalainen, S. U. Cho, F. Massel *et al.*, *Cavity optomechanics mediated by a quantum two-level system*, Nat. Commun. **6**, 6981 (2015).
- [287] F. Rouxinol, Y. Hao, F. Brito *et al.*, *Measurements of nanoresonator-qubit interactions in a hybrid quantum electromechanical system*, Nanotechnology **27**, 364003 (2016).
- [288] V. Singh, S. J. Bosman, B. H. Schneider *et al.*, *Optomechanical coupling between a multilayer graphene mechanical resonator and a superconducting microwave cavity*, Nat. Nanotechnol. **9**, 820–824 (2014).
- [289] B. H. Schneider, V. Singh, W. J. Venstra *et al.*, *Observation of decoherence in a carbon nanotube mechanical resonator*, Nat. Commun. **5**, 5819 (2014).
- [290] J. Moser, A. Eichler, J. Güttinger *et al.*, *Nanotube mechanical resonators with quality factors of up to 5 million*, Nat. Nanotechnol. **9**, 1007–1011 (2014).
- [291] N. Ares, T. Pei, A. Mavalankar *et al.*, *Resonant Optomechanics with a Vibrating Carbon Nanotube and a Radio-Frequency Cavity*, Phys. Rev. Lett. **117**, 170801 (2016).
- [292] P. Weber, J. Güttinger, I. Tsioutsios *et al.*, *Coupling Graphene Mechanical Resonators to Superconducting Microwave Cavities*, Nano Letters **14**, 2854–2860 (2014).
- [293] X. Song, M. Oksanen, J. Li *et al.*, *Graphene Optomechanics Realized at Microwave Frequencies*, Phys. Rev. Lett. **113**, 027404 (2014).
- [294] I. Kozinsky, H. W. C. Postma, I. Bargatin *et al.*, *Tuning nonlinearity, dynamic range, and frequency of nanomechanical resonators*, Appl. Phys. Lett. **88**, 253101 (2006).
- [295] M. D. LaHaye, *private communication* (2017).
- [296] F. Petruccione and H. Breuer, *The Theory of Open Quantum Systems*, Oxford University Press, Oxford, UK (2002).
- [297] H. W. C. Postma, I. Kozinsky, A. Husain *et al.*, *Dynamic range of nanotube- and nanowire-based electromechanical systems*, Appl. Phys. Lett. **86**, 223105 (2005).
- [298] S. Rips, I. Wilson-Rae and M. J. Hartmann, *Nonlinear nanomechanical resonators for quantum optoelectromechanics*, Phys. Rev. A **89**, 013854 (2014).
- [299] M. Agarwal, S. A. Chandorkar, R. N. Candler *et al.*, *Optimal drive condition for nonlinearity reduction in electrostatic microresonators*, Appl. Phys. Lett. **89**, 214105 (2006).
- [300] R. Khan, F. Massel and T. T. Heikkilä, *Tension-induced nonlinearities of flexural modes in nanomechanical resonators*, Phys. Rev. B **87**, 235406 (2013).

- [301] Y.-X. Liu, A. Miranowicz, Y. B. Gao *et al.*, *Qubit-induced phonon blockade as a signature of quantum behavior in nanomechanical resonators*, Phys. Rev. A **82**, 032101 (2010).
- [302] I. M. Pop, K. Geerlings, G. Catelani *et al.*, *Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles*, Nature **508**, 369–372 (2014).
- [303] J. J. Viennot, X. Ma and K. W. Lehnert, *Phonon-Number-Sensitive Electromechanics*, Phys. Rev. Lett. **121**, 183601 (2018).
- [304] S. Bravyi, D. Gosset and R. König, *Quantum advantage with shallow circuits*, Science **362**, 308–311 (2018).
- [305] V. Dunjko, Y. Ge and J. I. Cirac, *Computational Speedups Using Small Quantum Devices*, Phys. Rev. Lett. **121**, 250501 (2018).
- [306] F. Tacchino, C. Macchiavello, D. Gerace *et al.*, *An artificial neuron implemented on an actual quantum processor*, npj Quantum Inf. **5**, 26 (2019).
- [307] F. Tacchino, P. Barkoutsos, C. Macchiavello *et al.*, *Quantum implementation of an artificial feed-forward neural network*, arXiv:1912.12486 (2019).
- [308] G. Carleo, I. Cirac, K. Cranmer *et al.*, *Machine learning and the physical sciences*, arXiv:1903.10563 (2019).
- [309] M. Schuld, I. Sinayskiy and F. Petruccione, *An introduction to quantum machine learning*, Contemporary Physics **56**, 172–185 (2015).
- [310] A. Perdomo-Ortiz, M. Benedetti, J. Realpe-Gómez *et al.*, *Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers*, Quantum Sci. Technol. **3**, 030502 (2018).
- [311] M. Benedetti, E. Lloyd and S. Sack, *Parameterized quantum circuits as machine learning models*, arXiv:1906.07682 (2019).
- [312] Y. Nam, Y. Su and D. Maslov, *Approximate Quantum Fourier Transform with $O(n \log(n))$ T gates*, arXiv:1803.04933 (2018).
- [313] L. Wossnig, Z. Zhao and A. Prakash, *Quantum Linear System Algorithm for Dense Matrices*, Phys. Rev. Lett. **120**, 050502 (2018).
- [314] A. W. Harrow, A. Hassidim and S. Lloyd, *Quantum Algorithm for Linear Systems of Equations*, Phys. Rev. Lett. **103**, 150502 (2009).
- [315] V. Giovannetti, S. Lloyd and L. Maccone, *Quantum Random Access Memory*, Phys. Rev. Lett. **100**, 160501 (2008).
- [316] S. Lloyd, M. Mohseni and P. Rebentrost, *Quantum principal component analysis*, Nat. Phys. **10**, 631–633 (2014).
- [317] S. Lloyd, M. Mohseni and P. Rebentrost, *Quantum algorithms for supervised and unsupervised machine learning*, arXiv:1307.0411 (2013).
- [318] M. Schuld, M. Fingerhuth and F. Petruccione, *Implementing a distance-based classifier with a quantum interference circuit*, EPL (Europhysics Letters) **119**, 60002 (2017).
- [319] M. Schuld, A. Bocharov, K. Svore *et al.*, *Circuit-centric quantum classifiers*, arXiv:1804.00633 (2018).
- [320] E. Farhi and H. Neven, *Classification with Quantum Neural Networks on Near Term Processors*, arXiv:1802.06002 (2018).

- [321] V. Havlíček, A. D. Córcoles, K. Temme *et al.*, *Supervised learning with quantum-enhanced feature spaces*, *Nature* **567**, 209 (2019).
- [322] M. Schuld and N. Killoran, *Quantum Machine Learning in Feature Hilbert Spaces*, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [323] R. Ghobadi, J. S. Oberoi and E. Zahedinejhad, *The Power of One Qubit in Machine Learning*, arXiv:1905.01390 (2019).
- [324] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster *et al.*, *Data re-uploading for a universal quantum classifier*, arXiv:1907.02085 (2019).
- [325] A. W. Cross, G. Smith and J. A. Smolin, *Quantum learning robust against noise*, *Phys. Rev. A* **92**, 012327 (2015).
- [326] D. Ristè, M. P. da Silva, C. A. Ryan *et al.*, *Demonstration of quantum advantage in machine learning*, *npj Quantum Inf.* **3**, 16 (2017).
- [327] H. Buhrman, R. Cleve, J. Watrous *et al.*, *Quantum Fingerprinting*, *Phys. Rev. Lett.* **87**, 167902 (2001).
- [328] N. Wiebe, A. Kapoor and K. Svore, *Quantum Algorithms for Nearest-Neighbor Methods for Supervised and Unsupervised Learning*, arXiv:1401.2142 (2014).
- [329] P. Rebentrost, M. Mohseni and S. Lloyd, *Quantum Support Vector Machine for Big Data Classification*, *Phys. Rev. Lett.* **113**, 130503 (2014).
- [330] W. S. McCulloch and W. Pitts, *A logical calculus of the ideas immanent in nervous activity*, *The bulletin of mathematical biophysics* **5**, 115–133 (1943).
- [331] F. Rosenblatt, *The Perceptron: A perceiving and recognizing automaton*, Tech. Rep. 85-460-1, Cornell Aeronautical Laboratory, Inc. (1957).
- [332] D. E. Rumelhart, G. E. Hinton and R. J. Williams, *Learning representations by back-propagating errors*, *Nature* **323**, 533–536 (1986).
- [333] G. Cybenko, *Approximation by superpositions of a sigmoidal function*, *Math. Control Signals Syst.* **2**, 303–314 (1989).
- [334] K. Hornik, *Approximation capabilities of multilayer feedforward networks*, *Neural Networks* **4** (1991).
- [335] J. M. Zurada, *Introduction to Artificial Neural Systems*, West Group (1992).
- [336] R. Rojas, *Neural Networks: A Systematic Introduction*, Springer (1996).
- [337] J. J. Hopfield, *Neural networks and physical systems with emergent collective computational abilities*, *PNAS* **79**, 2554 (1982).
- [338] R. McEliece, E. Posner, E. Rodemich *et al.*, *The capacity of the Hopfield associative memory*, *IEEE Transactions on Information Theory* **33** (1987).
- [339] M. Schuld, I. Sinayskiy and F. Petruccione, *The quest for a Quantum Neural Network*, *Quantum Inf. Process.* **13**, 2567–2586 (2014).
- [340] P. Rebentrost, T. R. Bromley, C. Weedbrook *et al.*, *Quantum Hopfield neural network*, *Phys. Rev. A* **98**, 042308 (2018).
- [341] M. V. Altaisky, *Quantum neural network*, arXiv:quant-ph/0107012 (2001).
- [342] M. Schuld, I. Sinayskiy and F. Petruccione, *Simulating a perceptron on a quantum computer*, *Phys. Lett. A* **379**, 660–663 (2015).

- [343] W. Hu, *Towards a Real Quantum Neuron*, *Natural Science* **10**, 99–109 (2018).
- [344] E. Torrontegui and J. J. Garcia-Ripoll, *Unitary quantum perceptron as efficient universal approximator*, *EPL (Europhysics Letters)* **125** (2019).
- [345] Y. Cao, G. G. Guerreschi and A. Aspuru-Guzik, *Quantum Neuron: an elementary building block for machine learning on quantum computers*, arXiv:1711.11240 (2017).
- [346] B. Ricks and D. Ventura, *Training a Quantum Neural Network*, *Advances in Neural Information Processing Systems* **16**, 1–8 (2003).
- [347] N. Wiebe, A. Kapoor and K. M. Svore, *Quantum Perceptron Models*, arXiv:1602.04799 (2016).
- [348] Y. Liao, O. Dahlsten, D. Ebler *et al.*, *Quantum advantage in training binary neural networks*, arXiv:1810.12948 (2018).
- [349] K. H. Wan, O. Dahlsten, H. Kristjánsson *et al.*, *Quantum generalisation of feedforward neural networks*, *npj Quantum Inf.* **3**, 36 (2017).
- [350] G. Verdon, J. Pye and M. Broughton, *A Universal Training Algorithm for Quantum Deep Learning*, arXiv:1806.09729 (2018).
- [351] I. Goodfellow, J. Pouget-Abadie, M. Mirza *et al.*, *Generative Adversarial Nets*, in *Proceedings of the International Conference on Neural Information Processing Systems*, pp. 2672–2680 (2014).
- [352] J. Romero, J. P. Olson and A. Aspuru-Guzik, *Quantum autoencoders for efficient compression of quantum data*, *Quantum Sci. Technol.* **2**, 045001 (2017).
- [353] L. Lamata, U. Alvarez-Rodriguez, J. D. Martín-Guerrero *et al.*, *Quantum autoencoders via quantum adders with genetic algorithms*, *Quantum Sci. Technol.* **4**, 014007 (2018).
- [354] Y. Ding, L. Lamata, M. Sanz *et al.*, *Experimental Implementation of a Quantum Autoencoder via Quantum Adders*, *Adv. Quantum Technol.* **2**, 1800065 (2019).
- [355] D. Silver, A. Huang, C. J. Maddison *et al.*, *Mastering the game of Go with deep neural networks and tree search*, *Nature* **529**, 484–489 (2016).
- [356] V. Dunjko, J. M. Taylor and H. J. Briegel, *Quantum-Enhanced Machine Learning*, *Phys. Rev. Lett.* **117**, 130501 (2016).
- [357] L. Lamata, *Basic protocols in quantum reinforcement learning with superconducting circuits*, *Sci. Rep.* **7**, 1609 (2017).
- [358] T. Sriarunothai, S. Wölk, G. S. Giri *et al.*, *Speeding-up the decision making of a learning agent using an ion trap quantum processor*, *Quantum Sci. Technol.* **4**, 015014 (2018).
- [359] J. S. Otterbach, R. Manenti, N. Alidoust *et al.*, *Unsupervised Machine Learning on a Hybrid Quantum Computer*, arXiv:1712.05771 (2017).
- [360] S. H. Adachi and M. P. Henderson, *Application of Quantum Annealing to Training of Deep Neural Networks*, arXiv:1510.06356 (2015).
- [361] M. H. Amin, E. Andriyash, J. Rolfe *et al.*, *Quantum Boltzmann Machine*, *Phys. Rev. X* **8**, 021050 (2018).
- [362] M. Benedetti, J. Realpe-Gómez, R. Biswas *et al.*, *Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning*, *Phys. Rev. A* **94**, 022308 (2016).

- [363] L. G. Wright and P. L. McMahon, *The Capacity of Quantum Neural Networks*, arXiv:1908.01364 (2019).
- [364] S. Ghosh, A. Opala, M. Matuszewski *et al.*, *Quantum reservoir processing*, npj Quantum Inf. **5**, 35 (2019).
- [365] L. B. Kristensen, M. Degroote, P. Wittek *et al.*, *An Artificial Spiking Quantum Neuron*, arXiv:1907.06269 (2019).
- [366] J. Schmidhuber, *Deep learning in neural networks: An overview*, Neural Networks p. 33 (2015).
- [367] C. D. Schuman, T. E. Potok, R. M. Patton *et al.*, *A Survey of Neuromorphic Computing and Neural Networks in Hardware*, arXiv:1705.06963 (2017).
- [368] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza *et al.*, *A million spiking-neuron integrated circuit with a scalable communication network and interface*, Science **345**, 668 (2014).
- [369] M. Sanz, L. Lamata and E. Solano, *Quantum memristors in quantum photonics*, APL Photonics **3**, 080801 (2018).
- [370] F. Neukart and S.-A. Moraru, *On Quantum Computers and Artificial Neural Networks*, Signal Processing Research **2**, 12 (2013).
- [371] M. Rossi, M. Huber, D. Bruß *et al.*, *Quantum hypergraph states*, New J. Phys. **15**, 113022 (2013).
- [372] M. Ghio, D. Malpetti, M. Rossi *et al.*, *Multipartite entanglement detection for hypergraph states*, J. Phys. A: Math. Theor. **51**, 045302 (2018).
- [373] C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer (2006).

List of Publications

Scientific Journals

1. **F. Tacchino**, A. Chiesa, S. Carretta and D. Gerace, *Quantum Computers as Universal Quantum Simulators: State-of-the-Art and Perspectives*, Advanced Quantum Technologies, 1900052 (2019). DOI: 10.1002/qute.201900052. Pre-print at arXiv:1907.03505.
2. **F. Tacchino**, A. Succurro, O. Ebenhöf and D. Gerace, *Optimal efficiency of the Q-cycle mechanism around physiological temperatures from an open quantum systems approach*, Scientific Reports **9**, 16657 (2019). DOI: 10.1038/s41598-019-52842-x. Pre-print at arXiv:1807.02332.
3. **F. Tacchino**, C. Macchiavello, D. Gerace and D. Bajoni, *An Artificial Neuron Implemented on an Actual Quantum Processor*, npj Quantum Information **5**, 26 (2019). DOI: 10.1038/s41534-019-0140-4. Pre-print at arXiv:1811.02266.
4. A. Chiesa, **F. Tacchino**, M. Grossi, P. Santini, I. Tavernelli, D. Gerace and S. Carretta, *Quantum hardware simulating four-dimensional inelastic neutron scattering*, Nature Physics **15**, 455 (2019). DOI: 10.1038/s41567-019-0437-4. Pre-print at arXiv:1809.07974.
5. **F. Tacchino**, A. Chiesa, M. D. LaHaye, S. Carretta and D. Gerace, *Electromechanical Quantum Simulators*, Physical Review B **97**, 214302 (2018). DOI: 10.1103/PhysRevB.97.214302. Pre-print at arXiv:1711.00051.
6. **F. Tacchino**, A. Auffèves, M. F. Santos and D. Gerace, *Steady State Entanglement beyond Thermal Limits*, Physical Review Letters **120**, 063604 (2018). DOI: 10.1103/PhysRevLett.120.063604. Pre-print at arXiv:1710.01377.

Pre-print manuscripts

7. **F. Tacchino**, P. Barkoutsos, C. Macchiavello, I. Tavernelli, D. Gerace and D. Bajoni, *Quantum implementation of an artificial feed-forward neural network*, arXiv:1912.12486 (2019).

Conference Proceedings

8. **F. Tacchino**, A. Chiesa, M. D. LaHaye, I. Tavernelli, S. Carretta and D. Gerace, *Digital Quantum Simulations of Spin Models on Hybrid Platform and Near-Term Quantum Processors*, Proceedings **12**, 24 (2019). DOI: 10.3390/proceedings2019012024. Pre-print at arXiv:1902.04971.
9. **F. Tacchino**, M. Grossi, D. Gerace, A. Chiesa, P. Santini, S. Carretta, and I. Tavernelli, *Efficient Quantum Simulation of Dynamic Correlations on Superconducting Quantum Computers*, Conference on Lasers and Electro-Optics, OSA Technical Digest (Optical Society of America, 2019), paper JTh5C.3 (**postdeadline**). DOI: 10.1364/CLEO_AT.2019.JTh5C.3.

Acknowledgments

In the last proposition of his *Tractatus Logico-Philosophicus*, L. Wittgenstein suggests that “whereof one cannot speak, thereof one must be silent.” It is certainly difficult to rival with such complete, authoritative and liberating conclusion, let alone disagree with its prescription. For once, I will however try to formulate into words the immense sense of gratitude that I feel for those who contributed, in many different ways, towards the successful completion of this work.

First and foremost, my respect and admiration go to my supervisor, Prof. Dario Gerace, who tirelessly encouraged and assisted me over the last few years. A mentor and a friend, he guided me outside the PhD box and into the fascinating world of scientific research. I fully owe my experience as a graduate student, well above the average in terms of personal and professional growth, to his initiative and openness. I also gratefully acknowledge the support and the many opportunities offered by the University of Pavia and by the Graduate School in Physics, the latter under the effective direction of Prof. L. C. Andreani.

I warmly thank Prof. L. Lamata and Prof. S. Montangero for their fast, careful review and very positive assessment of this dissertation.

I has been an honor and a pleasure to interact with many talented scientists around the world. I would like to thank in particular Dr. A. Auffèves and Prof. M. F. Santos, who welcomed me in their scientific family and accompanied me towards my first publication; Prof. S. Carretta and Dr. A. Chiesa, masters of molecular magnetism who gave crucial contributions to the results presented in Chapters 2 and 3 and to whom I am indebted for most of my knowledge in digital quantum simulations; Prof. D. Bajoni and Prof. C. Macchiavello, who got me involved, with their inventiveness and profound knowledge of quantum information, in the engaging quest for quantum artificial neurons; Dr. I. Tavernelli and Dr. P. Barkoutsos, who opened the doors of the IBM Research laboratory in Zurich, sharing their valuable expertise and insight on quantum algorithms, software and real processors; M. Grossi, who first introduced me to the IBM Q community and with whom I enjoyed several discussions regarding the business side of quantum computing; Prof. M. D. LaHaye for inviting me to visit Syracuse and for providing some of the beautiful images appearing in Chapter 3; Prof. H. E. Türeci, together with his team in Princeton and Berkeley, for his kind hospitality, support and understanding; Prof. O. Ebenhöf and Dr. A. Succurro for their perseverance. I would also like to mention Prof. M. Liscidini, who often took time to share with me his views on science and academic life, and my fellow graduate students M. Passoni, M. Clementi, F. Garrisi and A. Barone for good conversations and innumerable coffees.

Last but not least, a heartfelt thank you to L. Guglielmi for the beautiful cover illustration.

This work is dedicated to my parents, for hard work, integrity and patience are best taught by example. I shall also thank my sister and grandparents, who took care of me from opposite sides.

The final lines, and the deepest love, are for Elisa. This is the last safe haven of my youth, and here I am, once again, asking for your trust, strength and resilience. We both know I will only be able to face the journey of life with you by my side: come what may, I will always be standing by yours.