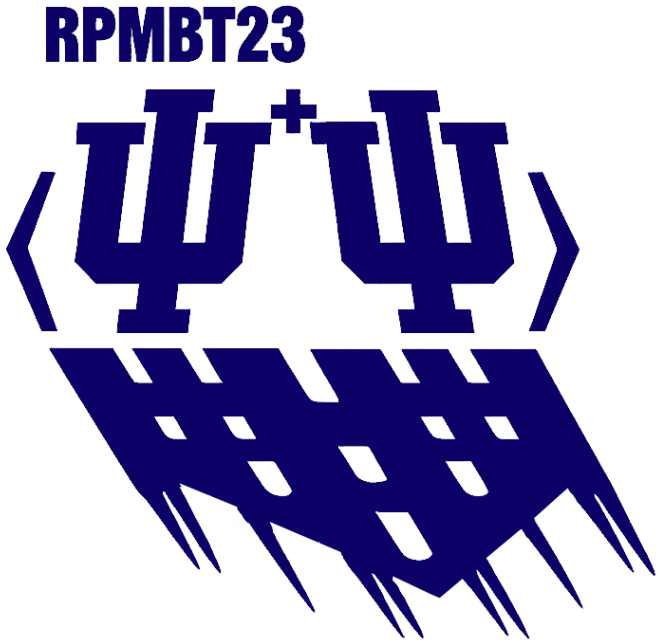

23rd International Conference on
Recent Progress in Many-Body Theories
Milan, Italy, 14–18 September 2026

SCHOOL ON QUANTUM SIMULATION

9-11 SEPTEMBER 2026

PRE–READING MATERIAL

DRAFT 2.0

Contents

Motivation and Practical Organization

1

Quantum Simulation of Many-Body Dynamics

This school is devoted to quantum simulation as a tool for studying real-time many-body dynamics, with a focus on problems that are intrinsically difficult for classical computers. The target audience is Master students, PhD students and early-stage researchers with a background in quantum many-body physics. Rather than providing a general overview of quantum computing or hardware architectures, the school concentrates on how to perform quantum simulations in practice, what physical regimes they address, and what their fundamental limitations are.

Motivation

Simulating the real-time dynamics of interacting quantum systems remains one of the central challenges of many-body physics. Classical methods are often limited by entanglement growth, exponential Hilbert space complexity, or sign problems. Quantum simulation offers a complementary approach by encoding and evolving quantum states directly on quantum degrees of freedom, avoiding the need for an explicit classical representation of the full many-body wavefunction. However, quantum simulation does not remove complexity; it redistributes it into circuit depth, entanglement, and algorithmic structure. Understanding this trade-off is a core theme of the school.

Hands-on access to quantum hardware

A practical component of the school will include reserved access to real quantum hardware. In addition to simulator-based exercises, participants will have the opportunity to run selected quantum simulation protocols on the IBM Quantum Platform, via classroom accounts, and the Turin Lagrange IQM Spark quantum computer, a superconducting five-qubit device intended for education and research. This access is meant to connect the conceptual material of the school with the concrete constraints of near-term quantum processors, while keeping the focus on many-body dynamics rather than on hardware engineering.

Conceptual framework

The school builds a coherent conceptual framework based on the following pillars:

- Locality and dynamics: How local interactions constrain the spread of correlations and structure quantum evolution.
- Digital quantum simulation: How many-body Hamiltonians are decomposed into local unitaries using Trotter-based algorithms, and how real-time dynamics is reconstructed step by step.
- Entanglement and complexity: Why entanglement growth and operator spreading impose fundamental limits on simulability, independently of hardware imperfections.

- Limits and opportunities: What quantum simulation can realistically achieve in the near term, and where it provides genuine conceptual or computational advantages.

Throughout, emphasis is placed on physical interpretation, not just algorithmic construction.

Prerequisites and recommended preliminary preparation

The school assumes familiarity with Quantum Mechanics and with the basic language of quantum circuits, qubits, gates, measurements, and simple circuit execution. These topics will not be reviewed systematically during the lectures. Participants who would benefit from a brief review of quantum mechanics are encouraged to read Sections 2.1 and 2.2 of Chapter 2 in [1] and/or Chapter 4 in [2]. Participants who wish to refresh the basic language of quantum circuits are encouraged to read Sections 1.1, 1.2 and 1.3 of Chapter 1 in [1] and Chapter 5 in [2].

Participants who do not yet feel comfortable with quantum computing basics, or with using a programming framework such as Qiskit, are strongly encouraged to complete some preliminary self-study before the school. The goal of this preparation is not to become experts in quantum algorithms or hardware, but to acquire enough operational familiarity to follow the discussion of digital quantum simulation protocols.

Recommended preliminary topics include:

- qubits, single-qubit gates, two-qubit gates, and measurements;
- the circuit model of quantum computation;
- basic notions of entanglement in the circuit language;
- construction and simulation of simple circuits using Qiskit;
- elementary use of statevector simulators, measurement sampling, and expectation values.

Suggested resources. For self-study, we recommend the following online courses from the *Courses* web page on the IBM Quantum Platform:

- *Basics of quantum information*
- *Use a quantum computer today*
- *Quantum computing in practice*
- On the web page of the QS School we will release also some notebooks (Participants are expected to study and understand these notebooks before the School begins) related to:
  - measurement of observables from a fixed given state and estimate of statistical errors
  - basic quantum gates and implementation of a quantum circuit to exponentiate Pauli operators
  - circuit identities and adaptation to connectivity constraints using the topology of the quantum hardware

Learning outcomes

After the school, participants should be able to:

- formulate a many-body dynamical problem as a quantum simulation task,
- construct and analyze Trotterized simulation protocols,

- understand the physical meaning and impact of Trotter errors,
- relate quantum simulation methods to classical many-body techniques,
- identify regimes where quantum simulation is promising and where it is not.

Perspective

Quantum simulation is presented neither as a universal solver nor as a speculative future technology, but as a developing scientific methodology, grounded in established many-body concepts and constrained by fundamental physical limits. The school aims to provide participants with a realistic, critical, and conceptually solid understanding of what it means to do quantum simulation of many-body dynamics today.

Three-Day School Program

Practical notes

- Each morning and afternoon block will include a break.
- Each hands-on session is organized around one or more notebooks
- Whenever possible, the same reference model is reused across the three days. A typical choice is a one-dimensional spin chain, such as the transverse-field Ising model.

Tentative time-table plan

T = theory lecture, E = exercise or hands-on session.

- First Day
  - **T1A**
    - Quantum simulation as a many-body workflow
    - Real-time evolution as the native simulation task
  - **E1A**
    - Implementation of quantum circuit for Trotter evolution for simple lattice models in 1D (Ising/Heisenberg)
    - Evaluation of local time-dependent observables, polarization/energy, using state-vector simulations
  - **T1B**
    - general Trotter-Suzuki decomposition
    - stochastic compilation techniques
    - calculation of two point correlators in time and relation to linear response
  - **E1B**
    - implementation of full protocol to compute two-point correlators for systems done in the morning session
    - control of the simulation error from Trotter
    - switch from state-vector to sampling simulations and control of shot noise
    - first simple run on quantum hardware
- Second Day
  - **T2A**
    - Ground states as initial states
    - Variational Quantum Eigensolver
  - **E2A**
    - Implementation of simple variational ground state preparation for the models implemented the previous day

- Progressive generalization of the simulation including first just shot noise and then a full noise emulator

T2B
- Real quantum computers and noise channels
- Error mitigation techniques: read-out mitigation, Richardson extrapolation

E2B
- Implementation of error mitigation strategies for ground state preparation
- test with noise emulators and finally quantum hardware

▶ Third Day

T3
- Advanced error mitigation techniques including TREX, Pauli-twirling, Noise renormalization and Dynamical-decoupling
- Linear Combination of Unitaries with application to eigenvalue estimation/eigenvector preparation using Hadamard test and Quantum Phase Estimation
- Advanced techniques for ground state preparation, time evolution and measurement including QITE, QSP and classical shadow tomography

E3
- extension of the simulations in the past few days to use the more advanced mitigation techniques seen in the morning and direct implementation on real hardware
- final challenge: how many Trotter steps can be done on the quantum hardware and still get overall error under control?

Concluding Discussion

Special emphasis is placed on the many-body physics perspective: quantum simulation is not a universal shortcut around complexity, but a way of representing, evolving, and probing quantum many-body systems using quantum degrees of freedom. The final discussion highlights when this representation is physically meaningful, where it may provide an advantage, and which limitations remain fundamental.

Pre–reading material purpose 2

Quantum Simulation of Many-Body Dynamics — Operational Foundations

Scope and philosophy

This document is intended as mandatory pre-reading for the school Quantum Simulation of Many-Body Dynamics. Its purpose is not to provide a general introduction to quantum mechanics or quantum computing (if you need, please, refer to the documents linked in the paragraph *Suggested resources*). Instead, it aims to establish a shared operational language for treating many-body quantum dynamics as an algorithmic problem. The central object of interest throughout this school is the real-time evolution of interacting quantum many-body systems, particularly in regimes where classical numerical methods face fundamental limitations [3–5]. The focus is therefore on dynamics, not ground states, and on out-of-equilibrium phenomena, not equilibrium thermodynamics.

Assumed background

The intended audience consists of Master students, PhD students and early-stage researchers in many-body physics. The following background is assumed:

- Non-relativistic quantum mechanics
- Second quantization
- Standard lattice models (e.g. Ising)
- Basic familiarity with entanglement and reduced density matrices

The following topics are not reviewed in this document and will not be reintroduced during the school:

- Basic quantum mechanics
- Introduction to quantum information and quantum computing
- Hardware architectures

Readers unfamiliar with these topics are encouraged to consult standard references before proceeding, e.g. the web links we have provided in the Executive Summary.

What this document is not

It is important to clarify the scope by stating explicitly what this document does not attempt to do.

This document is not:

- a general introduction to quantum computing,
- a hardware-focused guide to noisy intermediate-scale quantum (NISQ) devices,
- a survey of quantum advantage claims,

- a replacement for standard many-body textbooks.

Instead, it is designed to prepare the reader to actively construct and analyze quantum simulations of many-body dynamics, with an emphasis on conceptual clarity and physical interpretation.

Units, conventions, and notation

Throughout this document we adopt the following conventions:

- Natural units are used, with $\hbar = 1$.
- Qubits are identified with spin–1/2 degrees of freedom.
- Tensor products between subsystems are implicit when no ambiguity arises.
- Time evolution is assumed to be generated by time–independent Hamiltonians unless explicitly stated otherwise.

Many-body dynamics as a computational problem

The time evolution of a closed quantum system is formally simple. Given a Hamiltonian, $\hat{H}$, and an initial state, $|\Psi(0)\rangle$, the state at time t is

$$|\Psi(t)\rangle = e^{-i\hat{H}t}|\Psi(0)\rangle$$

Despite its compact form, this expression conceals the core difficulty of quantum many-body physics: the exponential growth of the Hilbert space with system size.

For a lattice system of N spin-1/2 degrees of freedom, the dimension of the Hilbert space scales as 2^N. Exact diagonalization therefore becomes impractical already at moderate values of N, even before considering time evolution [6].

Classical approximate methods mitigate this problem by exploiting structure:

- tensor-network methods rely on low entanglement,
- semiclassical methods rely on weak quantum correlations,
- Monte Carlo methods rely on probabilistic sampling.

However, real-time dynamics generically destroys (or quickly complicates) this structure.

Why real-time dynamics is hard classically

Out-of-equilibrium dynamics presents a unique challenge to classical simulation methods [7, 8]:

- Exact diagonalization is limited by exponential memory and time requirements.
- Time-dependent density matrix renormalization group (DMRG) / time-evolving block decimation (TEBD) methods rely on low entanglement and typically break down after a finite simulation time due to linear entanglement growth following a global quench [9–11].

- Quantum Monte Carlo approaches are extremely powerful in imaginary time and in sign-problem-free regimes, but generic real-time unitary dynamics is hard to access: in most formulations one encounters either a dynamical sign/phase problem or an ill-conditioned analytic continuation from imaginary-time data (with only limited exceptions and uncontrolled approximations) [7].

The key obstruction is not system size alone, but the growth of entanglement during time evolution. After a global quench, entanglement entropy typically grows linearly in time, leading to an exponential increase in the classical resources required to represent the state [8, 11].

This observation motivates the search for alternative simulation paradigms that do not rely on entanglement compression.

Quantum simulation as an algorithmic approach

Quantum simulation addresses this challenge by directly reproducing quantum dynamics using quantum degrees of freedom. Rather than approximating the state within a restricted classical representation, one constructs an algorithm whose execution naturally generates the target evolution.

In this approach:

- the Hilbert space is not stored, but realized physically,
- entanglement is not approximated, but generated naturally,
- time evolution is implemented through sequences of local operations.

From this perspective, quantum simulation is best understood as an *algorithmic representation* of the target dynamics. In digital approaches, one approximates the unitary $e^{-i\hat{H}t}$ by decomposing it into a sequence of elementary evolutions (gates) associated with local terms of the Hamiltonian (Trotter–Suzuki and related constructions). In variational approaches, one instead tracks the evolution by projecting it onto a parametrized manifold $|\Psi(\theta(t))\rangle$ and integrating effective equations of motion for the parameters $\theta(t)$ (variational quantum simulation / TDVP-type methods).

The remainder of this document develops this viewpoint systematically, starting from local Hamiltonians and their causal structure, and progressing toward digital and variational quantum simulation schemes.

Many–Body Quantum Dynamics as a Computational Problem

3

3.1 Time evolution and representation of quantum states

The dynamics of a closed quantum system is governed by the Schrödinger equation

$$i\hbar\frac{d}{dt}|\Psi(t)\rangle = \hat{H}|\Psi(t)\rangle$$

whose formal solution is

$$|\Psi(t)\rangle = e^{-i\hat{H}t}|\Psi(0)\rangle$$

For lattice many-body systems, the Hamiltonian, $\hat{H}$, is typically expressed as a sum of local or quasi-local terms,

$$\hat{H} = \sum_l \hat{h}_l$$

where each $\hat{h}_l$ acts non-trivially only on a finite subset of degrees of freedom.

The exponential operator $e^{-i\hat{H}t}$ defines a unitary transformation on the many–body Hilbert space. The central computational challenge lies not in writing down this expression, but in representing and manipulating the quantum state on which it acts.

For a system of N spin-1/2 degrees of freedom, the Hilbert space dimension scales as

$$\dim\mathcal{H} = 2^N \quad .$$

Any explicit representation of the state vector therefore requires a number of parameters that grows exponentially with system size.

3.2 Exact methods and their limitations

The most direct approach to quantum dynamics is exact diagonalization. By diagonalizing $\hat{H}$ one can in principle compute $e^{-i\hat{H}t}$ exactly and apply it to an arbitrary initial state.

In practice, this approach is limited to very small systems. Memory requirements scale as $\mathcal{O}(2^N)$, and computational costs scale even more unfavorably. As a result, exact diagonalization is typically restricted to $N \lesssim 20$ for spin systems (order-of-magnitude, depending on symmetries and sparsity), and even smaller sizes for fermionic models [6].

While exact diagonalization provides valuable benchmarks, it is not a viable tool for studying dynamics in the thermodynamic limit or at experimentally relevant system sizes.

3.3 Tensor-network methods and entanglement growth

Tensor-network approaches, such as matrix product states (MPS) and time-evolving block decimation (TEBD), overcome the exponential scaling of exact methods by exploiting low entanglement [8–10].

For one-dimensional systems in or near their ground state, the entanglement entropy of a subsystem typically obeys an area law, allowing an efficient representation of the state using a number of parameters that scales polynomially with system size [8, 12].

However, this advantage is fragile under time evolution. Following a global quench, entanglement entropy generically grows linearly in time, $S(t) \sim v_E t$, where v_E is an entanglement velocity [11].

As a consequence, the bond dimension required to faithfully represent the state typically increases exponentially with time (for fixed accuracy), leading to a rapid breakdown of tensor-network simulations. In practice, this limits time-dependent DMRG simulations to relatively short timescales, even in one-dimensional systems.

The breakdown is not due to numerical instability, but to a fundamental growth of quantum correlations.

3.4 Quantum Monte Carlo and real-time dynamics

Quantum Monte Carlo (QMC) methods provide a powerful framework for studying equilibrium properties of many-body systems by sampling imaginary-time path integrals [6].

Real-time dynamics, however, poses a fundamental challenge to QMC due to the sign problem [7]. The oscillatory nature of real-time amplitudes leads to severe cancellations in stochastic sampling, rendering the approach impractical for generic interacting systems.

While analytic continuation from imaginary time is sometimes employed, it is an ill-posed problem and does not provide reliable access to real-time dynamics in general[13].

As a result, standard (sampling-based) QMC is generally not considered a viable tool for direct simulation of generic real-time unitary dynamics in strongly correlated systems. That said, "Monte Carlo"-flavoured variational approaches—such as time-dependent variational Monte Carlo (tVMC), where one evolves a parametrized wavefunction and estimates the required expectation values stochastically—can access certain real-time regimes at the price of introducing an explicit variational bias[14, 15].

3.5 Ground-State Problems versus Real-Time Dynamics

3.5.1 Static properties and classical success stories

A large fraction of computational many-body physics is devoted to the study of ground-state properties and low-energy spectra of interacting quantum systems. For these problems, a wide range of powerful classical methods exists.

Examples include:

- density matrix renormalization group (DMRG),
- variational Monte Carlo (VMC),
- tensor-network methods in one and higher dimensions,
- exact quantum Monte Carlo methods in sign–problem–free regimes.

These methods exploit specific structural properties of ground states, such as low entanglement, locality of correlations, or probabilistic interpretability.

As a result, ground-state physics is often accessible to classical simulation even for large system sizes.

3.5.2 Why is dynamics fundamentally different?

Real-time dynamics is qualitatively different from static problems.

Time evolution does not preserve the special structure that makes ground states classically tractable. Even when starting from a low-entanglement initial state, unitary evolution under an interacting Hamiltonian typically generates highly entangled states with volume-law entanglement.

This distinction is crucial: the difficulty of simulating dynamics is not inherited from the difficulty of solving the static problem, but arises from the properties of unitary time evolution itself.

In particular, systems whose ground states are efficiently representable may nevertheless exhibit dynamics that are classically intractable after short times.

3.5.3 Imaginary time versus real time

Many classical algorithms effectively simulate imaginary-time evolution, either explicitly or implicitly.

Imaginary-time propagation,

$$|\Psi(\tau)\rangle = e^{-\hat{H}\tau}|\Psi(0)\rangle$$

acts as a projection onto low-energy states and suppresses high-energy components. This process reduces entanglement and stabilizes numerical representations.

Real-time evolution,

$$|\Psi(t)\rangle = e^{-i\hat{H}t}|\Psi(0)\rangle$$

by contrast, preserves the norm and explores the full unitary orbit of the state. No natural compression mechanism exists, and entanglement is generically amplified rather than suppressed.

This fundamental difference explains why methods successful in imaginary time often fail in real time.

3.5.4 Variational principles: static versus dynamical

Variational principles play different roles in static and dynamical settings.

In ground-state problems, variational methods aim to minimize the energy over a restricted class of states. The quality of the result depends on how well the variational manifold captures the true ground state.

In dynamical problems, variational principles—such as the time-dependent variational principle—approximate the trajectory of the system in Hilbert space rather than a single state.

This shift from optimizing a scalar quantity to approximating a trajectory reflects the increased complexity of real-time dynamics.

3.5.5 Implications for quantum simulation

The distinction between static and dynamical problems has direct implications for quantum simulation.

Quantum simulation is not primarily motivated by ground-state calculations, where classical methods are often competitive. Its natural domain is real-time evolution, where classical methods encounter fundamental obstacles.

This does not imply that quantum simulation will outperform classical methods in all dynamical regimes. However, it clarifies why dynamics is the central target of quantum simulation efforts.

3.6 Entanglement as the key obstruction

Across different classical approaches, a unifying obstruction emerges: entanglement growth.

The difficulty of simulating quantum dynamics is not solely determined by system size or interaction strength, but by the structure of quantum correlations generated during time evolution. Even systems with local interactions and simple initial states can rapidly develop highly entangled states that defy efficient classical representation.

From a computational perspective, entanglement plays a role analogous to chaos in classical systems: it signals the breakdown of compact descriptions and leads to exponential complexity.

This observation motivates the exploration of simulation strategies that do not rely on entanglement compression.

3.7 Quantum simulation: shifting the representation paradigm

Quantum simulation offers a fundamentally different approach to the problem of many-body dynamics. Rather than attempting to store and manipulate the full quantum state on a classical computer, one realizes the Hilbert space directly using quantum degrees of freedom.

In this paradigm:

- quantum states are represented physically rather than numerically,
- entanglement is generated naturally by the dynamics,
- the cost of simulation is measured in terms of quantum resources such as circuit depth or evolution time, rather than classical memory.

The task of simulating dynamics is thus reformulated as the problem of decomposing the target unitary evolution into a sequence of physically realizable operations.

3.8 Computational resources and scaling

While quantum simulation avoids the exponential memory overhead of classical methods, it is not free of constraints. The complexity of a quantum simulation is governed by different resource measures, including:

- the number of degrees of freedom required to represent the system,
- the depth of the quantum circuit or the duration of analog evolution,
- the locality of interactions and the resulting causal structure.

Understanding how these resources scale with system size and simulation time is essential for assessing the feasibility and limitations of quantum simulation approaches.

These questions will be addressed in the following sections, starting with the structure of local Hamiltonians and their implications for information propagation and circuit complexity.

Local Hamiltonians, Locality, and Causality

4

4.1 Structure of local many-body Hamiltonians

Most quantum many-body systems of physical interest are governed by local Hamiltonians, i.e. Hamiltonians that can be written as sums of terms acting non-trivially only on finite, spatially localized subsets of degrees of freedom. Formally, one writes

$$\hat{H} = \sum_l \hat{h}_l \quad ,$$

where each operator $\hat{h}_l$ acts on a bounded region of the lattice.

Locality should be read operationally: it means that the Hamiltonian can be decomposed into pieces that involve only a few nearby degrees of freedom at a time. This is not merely a modelling convenience but a statement about which couplings are physically available. In many platforms (solid-state, cold atoms, ions), interactions are either genuinely short-ranged or decay sufficiently fast with distance that a local description is accurate over the timescales of interest.

Locality is therefore a physical constraint arising from short–range interactions and the resulting effective finite-speed propagation of information in lattice systems (as formalized by Lieb–Robinson bounds). It is also the key structural property that makes quantum simulation possible in practice.

Typical examples include:

- nearest-neighbor spin interactions,
- finite-range hopping and interactions in fermionic models,
- gauge-invariant local plaquette and link terms in lattice gauge theories.

Although the individual terms $\hat{h}_l$ are simple, they generally do not commute with each other,

$$\left[\hat{h}_l, \hat{h}_{l'}\right] \neq 0$$

which makes the exact time evolution non–trivial.

4.2 Locality versus non–commutativity

The simultaneous presence of locality and non-commutativity is the defining feature of interacting quantum systems.

If all local terms commuted, the time evolution operator would factorize exactly:

$$e^{-i\hat{H}t} = \prod_l e^{-i\hat{h}_l t} \quad .$$

In such a case, quantum dynamics would be trivially simulable.

When does the evolution factorize?

If $[\hat{h}_l, \hat{h}_{l'}] = 0$ for all pairs, then all terms can be simultaneously diagonalized and one can rewrite

$$e^{-it\sum_l \hat{h}_l} = \prod_l e^{-it\hat{h}_l} .$$

When commutators do not vanish, the Baker–Campbell–Hausdorff structure implies additional nested-commutator terms, which is precisely the origin of Trotter errors in digital simulation.

Interacting systems, however, are characterized precisely by the fact that local terms overlap and fail to commute. This leads to:

- generation of entanglement,
- propagation of correlations,
- non-trivial dynamical behavior.

From the perspective of quantum simulation, the problem is therefore to approximate the global evolution generated by non-commuting local terms using sequences of local evolutions.

4.3 Causality and information propagation

Despite the absence of relativistic constraints in non-relativistic lattice systems, locality imposes strong restrictions on how information and correlations propagate in time.

This is formalized by the Lieb–Robinson bound, which states that for local Hamiltonians there exists an effective maximum velocity, v_{LR}, such that operators initially localized in one region of the lattice can influence distant regions only within a causal cone [16, 17].

Schematically, for two local operators A_x and B_y supported at positions x and y, one has

$$\|[A_x(t), B_y]\| \lesssim \exp\left\{\left[-\frac{|x-y| - v_{LR}t}{\xi}\right]\right\} \quad ,$$

where ξ is a characteristic length scale.

The Lieb–Robinson bound implies the existence of an emergent light cone in lattice systems, within which correlations can build up dynamically.

Why a light cone emerges (idea)

A useful starting point is the Heisenberg-evolved operator $A_x(t)$ and the commutator $C(t) = [A_x(t), B_y]$. Using $\frac{d}{dt}A_x(t) = i[\hat{H}, A_x(t)]$ one obtains

$$\frac{d}{dt}C(t) = i\left[\hat{H}, C(t)\right] + i\left[\left[\hat{H}, B_y\right], A_x(t)\right] .$$

Because $\hat{H} = \sum_l \hat{h}_l$ is local, the second term can be non-zero only when some $\hat{h}_l$ overlaps with the support of B_y. Iterating this argument generates a series of nested commutators: to connect x to y one needs a sequence of overlapping local terms, and the number of such

sequences is strongly constrained. Bounding the norm of each step yields an exponential suppression until $t \gtrsim |x-y|/v_{LR}$.

4.4 Physical interpretation of the Lieb–Robinson bound

The significance of the Lieb–Robinson bound goes beyond its formal statement. Physically, it implies that:

- correlations propagate at finite speed,
- distant regions remain approximately independent at short times,
- entanglement growth is constrained by locality.

These properties are essential for understanding both the limitations of classical simulations and the structure of quantum simulation algorithms.

In particular, the existence of a light cone implies that long-range correlations cannot be generated instantaneously, but require a sequence of local interactions over time.

4.5 Locality and circuit representations

The causal structure imposed by locality has a natural representation in terms of quantum circuits.

Time evolution under a local Hamiltonian can be viewed as a sequence of layers of local operations, each acting on neighboring degrees of freedom. In this picture:

- each local term $\hat{h}_l$ generates a local unitary evolution,
- non-commuting terms require temporal ordering,
- the depth of the circuit grows with simulation time.

The emergent light cone of the Lieb–Robinson bound corresponds to the spreading of operator support under successive layers of local gates.

This correspondence provides an intuitive bridge between many–body physics and quantum simulation: local Hamiltonian dynamics maps naturally onto local quantum circuits.

4.6 Example: 1D spin chains (two complementary decompositions)

To make these ideas concrete, consider a one-dimensional spin-1/2 chain. In practice, there are two complementary ways of decomposing the Hamiltonian for simulation purposes, and it is useful to keep them conceptually distinct:

- a *physics-motivated split* into simple operator families (e.g. interactions versus fields), which often leads to commuting blocks and clear gate syntheses;
- a *scheduling split* (odd/even bonds) used to parallelize two-site gates when adjacent terms do not commute.

Physics-motivated split (the transverse-field Ising model). For the transverse-field Ising chain one writes

$$\hat{H} = \hat{H}_{ZZ} + \hat{H}_X \quad , \qquad \hat{H}_{ZZ} = -J \sum_j \sigma_j^z \sigma_{j+1}^z \, , \qquad \hat{H}_X = -h \sum_j \sigma_j^x \, .$$

All terms within $\hat{H}_{ZZ}$ commute with each other, and all terms within $\hat{H}_X$ commute with each other. This is exactly the decomposition used later when introducing Trotter steps for the Ising example.

Odd/even scheduling (generic nearest-neighbour two-site interactions). More generally, if the Hamiltonian is written as a sum of nearest-neighbour two-site terms

$$\hat{H} = \sum_j \hat{h}_{j,j+1} \, ,$$

then terms on disjoint bonds commute, while adjacent-bond terms typically do not. In 1D this implies a natural two-colouring of the interaction graph:

$$\hat{H} = \hat{H}_{\text{odd}} + \hat{H}_{\text{even}} \quad , \qquad \hat{H}_{\text{odd}} = \sum_{j \,\text{odd}} \hat{h}_{j,j+1} \quad , \quad \hat{H}_{\text{even}} = \sum_{j \,\text{even}} \hat{h}_{j,j+1} \, .$$

Operationally, one applies a layer of commuting two-site gates on odd bonds followed by a layer on even bonds. This is the basic mechanism behind finite-depth circuit layers in 1D and is the scheduling ingredient that complements (rather than replaces) a physics-motivated split such as $\hat{H}_{ZZ} + \hat{H}_X$.

4.7 Fermionic systems and apparent non-locality

For fermionic models, locality is more subtle. While fermionic Hamiltonians are local in terms of fermionic operators, their representation in terms of qubits typically involves non-local mappings, such as the Jordan–Wigner transformation (or alternative encodings such as Bravyi–Kitaev) [18, 19].

Under such mappings, local fermionic operators may acquire non-local string operators. Importantly, this non-locality is representational, not physical.

The underlying dynamics remains governed by a local Hamiltonian and therefore obeys Lieb–Robinson bounds. However, the apparent non-locality introduces additional overhead in quantum simulation algorithms and must be accounted for carefully.

4.8 Locality as an algorithmic resource

From the perspective of quantum simulation, locality is not merely a physical constraint but an algorithmic resource.

Locality implies:

- finite-depth circuit approximations at short times,
- bounded growth of operator support,

- structured decompositions of time evolution.

These properties enable the construction of quantum simulation algorithms whose complexity scales polynomially with system size for fixed simulation times [4, 20].

At the same time, locality also implies fundamental limits: simulating long-time dynamics or generating long-range correlations necessarily requires deep circuits.

Understanding locality and causality is therefore essential for assessing both the power and the limitations of quantum simulation.

In the next Chapter, we will exploit the structure discussed here to construct explicit algorithms for simulating time evolution. In particular, we will introduce digital quantum simulation schemes based on Trotter–Suzuki decompositions, which approximate the global evolution operator using sequences of local unitaries.

These constructions make the abstract connection between locality, causality, and circuit complexity explicit, and provide the practical foundation for quantum simulation of many-body dynamics.

4.9 Entanglement Growth, Light Cones, and Circuit Complexity

4.9.1 Entanglement as a dynamical quantity

Entanglement is not a static property of a quantum many-body system, but a dynamical quantity generated by time evolution. Even when starting from simple product states, interacting local Hamiltonians generically produce entanglement at a finite rate.

For a bipartition of the system into subsystems A and B, the entanglement entropy

$$S_A(t) = -\operatorname{Tr}\left[\hat{\rho}_A(t)\log\hat{\rho}_A(t)\right]$$

typically grows linearly in time after a global quench [8, 11],

$$S_A(t) \sim v_E t \quad ,$$

where v_E is the entanglement production rate.

This behavior is observed in a wide range of interacting systems and reflects the fundamental role of interactions in spreading quantum correlations.

4.9.2 Entanglement growth and classical simulability

The linear growth of entanglement has direct computational consequences.

Classical simulation methods that rely on compressed representations of quantum states —such as matrix product states— require a bond dimension that grows exponentially with the entanglement entropy. As a result, the computational cost of classical time evolution increases exponentially with time, even for one-dimensional systems.

Importantly, this breakdown is not algorithm-dependent, but structural: no classical representation based on low–entanglement manifolds can efficiently capture volume–law entangled states.

This observation explains why real–time many–body dynamics remains one of the most challenging problems in computational physics.

4.9.3 Light cones and the structure of correlations

Entanglement growth is constrained by locality. As discussed in Section 2.3, the Lieb–Robinson bound implies the existence of an effective light cone within which correlations can propagate.

For local observables $\hat{O}_x$ and $\hat{O}_y$ separated by a distance $|x-y|$, correlations remain exponentially suppressed until times $t \gtrsim |x-y|/v_{LR}$.

This causal structure has two important implications:

1. Entanglement spreads locally, via successive interactions between neighboring degrees of freedom.
2. Long-range correlations require finite time to develop.

Thus, entanglement growth and light-cone propagation are two manifestations of the same underlying locality constraint.

4.9.4 From light cones to circuits

The emergence of light cones suggests a natural circuit representation of time evolution.

Consider a local Hamiltonian acting on a lattice. Time evolution over a small interval Δt can be approximated by a layer of local operations acting on neighboring sites. Repeating this process generates a quantum circuit with finite-range gates, whose depth grows linearly with the simulated time.

In this picture:

- the circuit depth controls how far information can propagate,
- each circuit layer expands the effective light cone by a finite amount,
- entanglement growth is directly tied to circuit depth.

This establishes a concrete relationship between physical time evolution and circuit complexity.

4.9.5 Circuit depth as a measure of dynamical complexity

From the perspective of quantum simulation, circuit depth emerges as a natural measure of the complexity of many-body dynamics.

Short-time dynamics corresponds to shallow circuits, which generate limited entanglement and correlations. Long-time dynamics requires deep circuits capable of generating volume–law entangled states.

This perspective clarifies an important point: quantum simulation does not eliminate dynamical complexity, but provides a framework in which that complexity can be generated and manipulated efficiently.

The difficulty of simulating long–time dynamics is therefore not an artifact of classical algorithms, but a reflection of the intrinsic complexity of quantum many-body evolution.

4.9.6 Why does quantum simulation target dynamics?

The considerations above explain why quantum simulation is particularly well-suited to problems involving real-time dynamics.

While classical methods struggle once entanglement grows beyond a certain threshold, a quantum simulator can, in principle, generate and sustain highly entangled states as part of its natural evolution.

This does not imply that all dynamical problems are easy on a quantum simulator. However, it suggests that quantum simulation aligns naturally with the structure of many-body dynamics, rather than fighting against it.

At this point, the ingredients are in place to introduce concrete quantum simulation algorithms.

- Local Hamiltonians define the structure of interactions.
- Locality imposes causal constraints.
- Entanglement growth sets the scale of dynamical complexity.
- Circuit depth provides a natural language to quantify resources.

In the next chapter, we will make this transition explicit by constructing digital quantum simulation schemes that approximate time evolution using sequences of local unitaries, starting from Trotter–Suzuki decompositions.

Digital Quantum Simulation of Many-Body Dynamics

5

5.1 Objective of digital quantum simulation

The goal of digital quantum simulation is to approximate the real-time evolution generated by a many-body Hamiltonian,

$$\hat{U}(t) = e^{-i\hat{H}t} \quad ,$$

using a sequence of elementary unitary operations associated with local terms of the Hamiltonian.

Unlike classical numerical integration of differential equations, digital quantum simulation does not attempt to compute the state explicitly. Instead, it constructs a quantum algorithm whose execution reproduces the desired unitary evolution.

The central problem is therefore algorithmic:

How can one decompose the global unitary $e^{-i\hat{H}t}$ into a sequence of local unitaries that can be applied sequentially?

5.2 Decomposition of the Hamiltonian

We assume that the Hamiltonian can be written as a sum of local terms,

$$\hat{H} = \sum_l \hat{h}_l \quad ,$$

In general, these terms do not commute. As a result,

$$e^{-i\hat{H}t} \neq \prod_l e^{-i\hat{h}_l t} \quad .$$

Digital quantum simulation proceeds by approximating the global evolution as a product of local evolutions over short time steps.

5.3 First–order Trotter decomposition

The simplest approximation is the first-order Trotter (Lie–Trotter) decomposition [21]. For a small time step Δt, one writes

$$e^{-i\hat{H}\Delta t} = e^{-i\sum_l \hat{h}_l \Delta t} \simeq \prod_l e^{-i\hat{h}_l \Delta t} \quad .$$

By repeating this step $n = t/\Delta t$ times, one approximates the full evolution,

$$e^{-i\hat{H}t} \simeq \left(\prod_l e^{-i\hat{h}_l \Delta t} \right)^n \quad .$$

This approximation becomes exact in the limit $\Delta t \to 0$.

5.4 Physical interpretation of Trotterization

It is important to emphasize that Trotterization is not a numerical integration scheme in the classical sense.

The Trotter approximation replaces the true dynamics generated by the full Hamiltonian with a stroboscopic dynamics, in which different interaction terms act sequentially rather than simultaneously.

Where does the Trotter error come from?

Consider the simplest non-trivial case $\hat{H} = \hat{A}+\hat{B}$. The Lie–Trotter step replaces $e^{-i(\hat{A}+\hat{B})\Delta t}$ with $e^{-i\hat{A}\Delta t}e^{-i\hat{B}\Delta t}$. Using the Baker–Campbell–Hausdorff expansion,

$$e^X e^Y = \exp\left(X + Y + \tfrac{1}{2}[X,Y] + \cdots\right),$$

with $X = -i\hat{A}\Delta t$ and $Y = -i\hat{B}\Delta t$, one finds an extra term $\propto \Delta t^2[\hat{A},\hat{B}]$. This is why the approximation becomes exact only as $\Delta t \to 0$ and why non-commutativity is the fundamental source of digital-simulation error.

The error arises because non-commuting terms generate additional dynamics when applied in sequence. To leading order,

$$e^{-i\hat{H}\Delta t} \simeq \prod_l e^{-i\hat{h}_l\Delta t} \exp\left\{\left(-\frac{\Delta t^2}{2}\sum_{l<l'}\left[\hat{h}_l,\hat{h}_{l'}\right] + \mathcal{O}(\Delta t^3)\right)\right\} \quad .$$

Thus, Trotter errors reflect the physical non-commutativity of interactions, not numerical inaccuracies.

5.5 Second–order (symmetric) Trotter decomposition

Higher accuracy can be achieved using symmetric decompositions [22]. For a Hamiltonian split into two non-commuting parts,

$$\hat{H} = \hat{H}_A + \hat{H}_B \quad ,$$

the second-order Trotter decomposition reads

$$e^{-i\hat{H}\Delta t} \simeq e^{-i\hat{H}_A\Delta t/2}e^{-i\hat{H}_B\Delta t}e^{-i\hat{H}_A\Delta t/2} \quad .$$

This approximation has an error of order $\mathcal{O}(\Delta t^3)$ per step and preserves time-reversal symmetry.

Why the symmetric step is more accurate?

The key point is that the symmetric product $e^{-i\hat{A}\Delta t/2}e^{-i\hat{B}\Delta t}e^{-i\hat{A}\Delta t/2}$ is invariant under $\Delta t \mapsto -\Delta t$ up to inversion. As a consequence, all even powers of Δt in the effective generator cancel, and the leading correction starts at order Δt^3.

Second-order decompositions are widely used in practice, as they significantly reduce Trotter errors at the cost of a modest increase in circuit depth.

5.6 Circuit depth and resource scaling

Each local exponential $e^{-i\hat{h}_l \Delta t}$ corresponds to a local quantum gate acting on the degrees of freedom involved in $\hat{h}_l$.

The total circuit depth required to simulate time t scales as

$$\text{depth} \sim \frac{t}{\Delta t} \times \text{number of layers per step}$$

This scaling reflects a fundamental trade-off:

- smaller Δt reduces Trotter error,
- smaller Δt increases circuit depth.

Thus, digital quantum simulation does not remove dynamical complexity, but repackages it into circuit depth. In practice, the relevant quantity is not the raw number of local exponentials but the number of *layers* once commuting terms are grouped (e.g. even/odd bonds in 1D). This is why locality and parallelization strongly affect feasibility on real devices.

5.7 Example: transverse–field Ising chain

To make the construction explicit, consider the one-dimensional transverse-field Ising model,

$$\hat{H} = -J \sum_j \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z - h \sum_j \hat{\sigma}_j^x$$

The Hamiltonian naturally decomposes into two parts,

$$\hat{H} = \hat{H}_{ZZ} + \hat{H}_X \quad ,$$

with

$$\hat{H}_{ZZ} = -J \sum_j \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z \quad , \quad \hat{H}_X = -h \sum_j \hat{\sigma}_j^x$$

All terms within each part commute. The second–order Trotter step reads

$$\hat{U}(\Delta t) \simeq e^{-i\hat{H}_X \Delta t/2} e^{-i\hat{H}_{ZZ} \Delta t} e^{-i\hat{H}_X \Delta t/2} \quad .$$

Each factor corresponds to a layer of commuting local operations and can be implemented in parallel.

From Hamiltonian terms to gates (Ising example)

For the transverse-field Ising chain, the single-site term generates a simple rotation $e^{+ih\Delta t\, \sigma^x} = R_x(-2h\Delta t)$ (up to conventions), while the two-site interaction generates a controlled phase. Concretely, for a single bond one can implement $e^{+iJ\Delta t\, \sigma^z \otimes \sigma^z}$ with a short two-qubit pattern such as CNOT–$R_z(2J\Delta t)$–CNOT. This illustrates the core workflow of digital simulation: translate each local exponential into a

small gate block, then schedule commuting blocks in parallel layers.

5.8 Quench dynamics and observables

A typical quantum simulation protocol involves:

- preparing an initial state $|\Psi(0)\rangle$,
- applying the Trotterized evolution operator,
- measuring observables at discrete times.

For the Ising chain, relevant observables include:

- local magnetization $\langle \hat{\sigma}_j^z(t) \rangle$
- equal–time correlation functions,
- domain–wall dynamics.

The quantum simulation reproduces the unitary dynamics directly, without constructing the full wavefunction.

5.9 Trotter error versus physical effects

An important conceptual point is the distinction between:

- genuine physical effects of the Hamiltonian,
- artifacts introduced by finite Trotter step size.

In particular, finite-step Trotter errors may:

- break integrability,
- introduce effective higher-order interactions,
- modify long–time behavior.

From a physical perspective, a finite-step Trotter circuit is a *stroboscopic* (periodically driven) dynamics: after each step Δt one applies the same sequence of non-commuting gates, so the evolution is generated by a one-step unitary $\hat{U}_{\Delta t}$ repeated many times. It is therefore natural to define an *effective (Floquet) Hamiltonian* $\hat{H}_F$ through

$$\hat{U}_{\Delta t} = e^{-i\hat{H}_F\,\Delta t} \quad .$$

For sufficiently small Δt, one has $\hat{H}_F = \hat{H} + \mathcal{O}(\Delta t)$ (more precisely, $\hat{H}_F = \hat{H} + \mathcal{O}(\Delta t^p)$ for a pth-order Trotter formula), so the stroboscopic dynamics agrees with the target Hamiltonian evolution up to controlled corrections at short times [23].

Understanding this effective description is essential when interpreting simulation results.

5.10 Beyond nearest-neighbor spin models

The same framework applies to more complex systems:

- XXZ chains require additional two-site terms but retain a similar structure.
- Fermionic models, such as the Hubbard chain, require mappings to qubits and introduce non-local strings.

- Gauge theories involve additional constraints and multi-body terms.

In all cases, the core idea remains unchanged: decompose the Hamiltonian into local terms, approximate the global evolution by sequential local evolutions, and analyze the resulting resource requirements.

5.11 Limitations of digital quantum simulation

Digital quantum simulation has intrinsic limitations:

- long-time dynamics requires deep circuits,
- Trotter errors accumulate over time,
- chaotic dynamics amplifies small errors.

These limitations are not artifacts of current technology, but reflect fundamental properties of many-body quantum dynamics.

As a result, digital quantum simulation is most effective for:

- short- to intermediate-time dynamics,
- local observables,
- controlled quench protocols.

5.12 Outlook: toward variational and hybrid approaches

The limitations of Trotter-based simulation motivate alternative approaches, such as variational quantum simulation, which aim to approximate dynamics within restricted manifolds of quantum states.

These methods trade exactness for reduced circuit depth and will be discussed in later sections.

At this point, however, the essential framework of digital quantum simulation is complete: a physically motivated, algorithmic construction that directly targets the hardest regime of many-body physics—real-time quantum dynamics.

Entanglement, Circuit Complexity, and Fundamental Limits

6

This chapter provides a qualitative bridge between the "how" of digital simulation (Trotterization, gate synthesis) and the "how far" (which times, observables, and regimes are realistically accessible). The organizing idea is that many-body dynamics is hard precisely because it generates entanglement and spreads information, and these two features translate directly into circuit depth and sensitivity to imperfections.

6.1 Entanglement as a computational resource

In the context of quantum simulation, entanglement plays a dual role. On the one hand, it is the primary obstruction to efficient classical simulation. On the other hand, it is a resource that quantum simulators are designed to generate and manipulate.

This duality is central to understanding both the promise and the limitations of quantum simulation. Entanglement is not an optional feature of many-body dynamics, but an unavoidable consequence of local interactions and unitary evolution [8].

From an algorithmic perspective, the key question is not whether entanglement is present, but how much entanglement must be generated to faithfully reproduce the target dynamics.

Entanglement entropy and why it tends to grow after a quench

Given a bipartition $A \cup \bar{A}$ of a pure state $|\Psi\rangle$, the reduced density matrix is $\rho_A = \mathrm{Tr}_{\bar{A}}|\Psi\rangle\langle\Psi|$ and the von Neumann entanglement entropy is

$$S_A = -\mathrm{Tr}(\rho_A \log \rho_A) \quad .$$

For 1D gapped ground states, S_A typically obeys an area law (it saturates with subsystem size), which is the reason why matrix-product-state methods can be efficient.

After a global quench, local interactions create entangled quasi-particle pairs that propagate in opposite directions; as pairs are shared across the cut, S_A increases roughly linearly in time until finite-size effects or saturation set in. This same linear-in-time entanglement growth is the mechanism behind the growth of classical simulation cost, and it is also what forces digital simulations toward deeper circuits at longer times.

6.2 Circuit complexity and dynamical evolution

Digital quantum simulation represents time evolution as a quantum circuit composed of local gates. The depth of this circuit provides a natural measure of the complexity of the simulated dynamics.

For local Hamiltonians, circuit depth grows linearly with simulation time. This reflects the finite speed at which correlations propagate, as dictated by locality and Lieb–Robinson bounds [16, 17].

Locality ⇒ a light cone (Lieb–Robinson intuition)

For local Hamiltonians one can bound (in operator norm) the growth of commutators between initially separated operators. A common form is

$$\|[\hat{O}_x(t), \hat{O}_y]\| \leq c\,\|\hat{O}_x\|\,\|\hat{O}_y\|\,\exp\left[-\frac{d(x,y)-vt}{\xi}\right],$$

where $d(x,y)$ is the lattice distance, v is a Lieb–Robinson velocity, and c, ξ are model-dependent constants.

Operationally, this means that to create correlations (or significant commutators) across a distance r one needs a time $t \gtrsim r/v$, and therefore a circuit with depth scaling at least like $\mathcal{O}(r)$ (up to parallelization details). This is the conceptual reason why "long time" and "long distance" are synonymous with "deep circuits" in local many-body dynamics.

In this framework:

- short-time dynamics corresponds to shallow circuits,
- long-time dynamics requires deep circuits,
- highly entangled states correspond to circuits of large depth.

Circuit depth thus plays a role analogous to system size in classical simulations.

6.3 Operator growth and scrambling

An alternative perspective on dynamical complexity comes from the study of operator growth.

Under Heisenberg time evolution, a local operator $\hat{O}_x$ evolves into a sum of increasingly non-local operators,

$$\hat{O}_x(t) = e^{i\hat{H}t}\hat{O}_x e^{-i\hat{H}t}$$

The support of $\hat{O}_x(t)$ spreads within a light cone, and its operator complexity increases over time. This phenomenon is closely related to information scrambling and quantum chaos, often quantified via out-of-time-order correlators (OTOCs) [24, 25].

A standard OTOC and the butterfly effect

A widely used diagnostic of scrambling is the (squared) commutator between two simple operators, for instance local Pauli operators $\hat{W}_x$ and $\hat{V}_y$,

$$C_{W,V}(t) = -\langle[\hat{W}_x(t), \hat{V}_y]^2\rangle \quad .$$

If $\hat{W}_x(t)$ remains local near x, then for y far away the commutator is (approximately) zero; as $\hat{W}_x(t)$ grows, the commutator becomes non-zero

when the operator front reaches y. This motivates defining a butterfly velocity as the speed at which the region of non-commutativity expands.

In chaotic systems, $C_{W,V}(t)$ can display a rapid growth regime (often described as exponential over an intermediate window), capturing the intuition that a small local perturbation can spread and affect distant degrees of freedom.

From the standpoint of quantum simulation, operator growth implies that faithfully reproducing dynamics requires circuits capable of generating increasingly non-local correlations.

6.4 Chaotic dynamics and sensitivity to errors

In chaotic many-body systems, small perturbations in the Hamiltonian or initial state can lead to rapidly diverging trajectories in operator space (the quantum "butterfly effect") [24, 25].

This sensitivity has important consequences:

- small Trotter errors can grow rapidly in time,
- long-time dynamics becomes increasingly difficult to approximate,
- precise reproduction of late-time behavior may be infeasible.

These effects are not unique to quantum simulation, but reflect intrinsic properties of chaotic quantum dynamics.

Why small gate/Trotter errors can become big at late times?

A convenient way to think about a digital-simulation error is as a perturbation of the target evolution. If the implemented circuit corresponds to evolution under an effective Hamiltonian $\hat{H} + \delta\hat{H}$ (for instance from finite-step Trotterization or coherent gate miscalibrations), then at time t the error in an observable $\hat{O}$ can be related to how strongly $\hat{O}(t)$ becomes non-local.

Heuristically, when $\hat{O}(t)$ has grown to have support on many sites, it can have a large overlap with the perturbation $\delta\hat{H}$, and the difference between the ideal and perturbed dynamics becomes easier to detect. In chaotic systems, this intuition is formalized by the same operator-growth diagnostics used for scrambling (e.g. OTOCs), which can predict rapid sensitivity once the operator front has spread over the relevant region.

As a result, quantum simulation should be understood as a controlled approximation, whose validity depends on the timescale and observables of interest.

6.5 Local observables versus global properties

Not all physical quantities are equally sensitive to dynamical complexity.

Local observables often converge more rapidly and are less sensitive to long-range correlations than global quantities. As a result, quantum

simulation can provide meaningful access to local physics even when the global state is highly entangled.

This distinction motivates a focus on:

- local expectation values,
- correlation functions at finite distances,
- dynamical response functions.

Global properties, such as full-state fidelity or entanglement entropy at late times, are generally more demanding and may lie beyond practical reach.

6.6 Limits imposed by locality

Locality imposes both enabling constraints and fundamental limits on quantum simulation.

On the enabling side, locality ensures that:

- correlations spread at finite speed,
- short-time dynamics is accessible with shallow circuits.

On the limiting side, locality implies that:

- generating long-range correlations requires deep circuits,
- simulating dynamics beyond the light-cone timescale is inherently costly.

These limits are not artifacts of digital simulation, but reflect the structure of the underlying physical system.

6.7 No free lunch: quantum simulation versus classical methods

It is important to emphasize that quantum simulation does not uniformly outperform classical methods.

In regimes of low entanglement or short correlation lengths, classical approaches may remain more efficient. Quantum simulation becomes advantageous primarily in regimes where classical methods fail due to entanglement growth or sign problems.

Thus, the appropriate question is not whether quantum simulation is superior in general, but for which problems and timescales it provides a meaningful advantage.

6.8 Implications for algorithm design

The considerations above have direct implications for the design of quantum simulation algorithms:

- target observables should be chosen carefully,
- simulation times should be matched to achievable circuit depth,
- approximation schemes should be evaluated in terms of physical, not purely mathematical, accuracy.

What Quantum Simulation Can — and Cannot — Do

7

7.1 Beyond hype: a realistic perspective

Quantum simulation is often presented as a universal solution to hard quantum many-body problems. This narrative is misleading.

Quantum simulators do not eliminate complexity; they redistribute it. What changes is not the fundamental difficulty of the problem, but the computational resources used to address it.

A realistic understanding of what quantum simulation can and cannot do is essential for meaningful scientific progress.

7.2 Problems that quantum simulation targets naturally

Quantum simulation is naturally suited to problems where the "native" language is real-time unitary dynamics and where classical descriptions struggle—for instance because entanglement grows rapidly, because Monte Carlo formulations suffer from sign problems [7], or because the observables of interest are intrinsically quantum.

Concrete examples are global and local quenches, transport in interacting systems, non-equilibrium phases of matter, and dynamical diagnostics of scrambling such as out-of-time-order correlators [24, 25].

In these regimes, quantum simulators offer a direct representation of the physics, rather than an indirect numerical approximation.

7.3 Timescales and observables

The practical reach of quantum simulation is constrained by achievable circuit depth and error accumulation. As a result, the most reliable results typically concern short to intermediate times, and observables that are local (or few-body) such as correlation functions at finite distance and dynamical response functions.

Late-time behavior, thermalization, and fine-grained spectral properties are significantly more challenging and should be approached with caution.

This limitation is not a technological shortcoming, but reflects intrinsic properties of many-body dynamics.

7.4 Approximation is unavoidable

All quantum simulation methods rely on approximations. In digital approaches, Trotterization approximates the target unitary; in variational approaches, one approximates the state trajectory by restricting it to an ansatz manifold; in analog simulation, the implemented device Hamiltonian approximates the intended model.

The key question is not whether an approximation is present, but whether it is controlled and physically meaningful.

Quantum simulation should be understood as a controlled modeling framework, not as a black-box solver.

7.5 Comparison with classical many-body methods

Classical many-body techniques remain indispensable.

Tensor networks, quantum Monte Carlo, and semiclassical approaches each excel in specific regimes. Quantum simulation complements these methods by targeting regimes where classical representations break down.

A productive research strategy therefore combines classical simulations (for validation and benchmarking) with quantum simulations (to explore beyond classical reach). The two approaches should be viewed as synergistic, not competitive.

7.6 The role of NISQ-era devices

Near-term quantum devices operate in a regime characterized by noise, finite coherence times, and limited connectivity [26].

Within these constraints, meaningful quantum simulation requires carefully chosen models, robust observables, and error-aware algorithm design. The NISQ era favors qualitative insights and proof-of-principle studies, rather than high-precision quantitative predictions.

7.7 Conceptual payoffs

Even when practical advantages are limited, quantum simulation offers significant conceptual benefits: it suggests new ways of thinking about dynamics and entanglement, provides operational definitions of complexity, and gives experimental access to otherwise abstract quantities.

These insights enrich many-body theory independently of computational speedups.

7.8 Open directions and research opportunities

Several open directions are particularly promising, including hybrid quantum–classical simulation schemes, adaptive variational manifolds, error-mitigated digital simulation [27, 28], and deeper connections between scrambling, chaos, and complexity.

These areas provide fertile ground for PhD-level research and interdisciplinary collaboration.

7.9 Final perspective

Quantum simulation is neither a miracle nor a dead end.

It is a developing scientific tool, grounded in well-understood physical principles and constrained by fundamental limits. Its power lies in addressing the right questions with the right expectations.

For many-body physics, quantum simulation offers not an escape from complexity, but a new way to engage with it.

Conclusions

This document has presented quantum simulation as a framework for addressing one of the central challenges of many-body physics: the study of real-time quantum dynamics in regimes where classical approaches become ineffective [3–5].

Rather than treating quantum simulation as a monolithic concept, we have emphasized its internal structure, limitations, and conceptual foundations. The focus has been on digital and variational approaches, which recast time evolution as an algorithmic process implemented through quantum circuits or constrained variational manifolds [21, 22, 29, 30].

Throughout, a recurring theme has been that quantum simulation does not eliminate complexity, but reshapes it. Locality, entanglement growth, circuit depth, and operator spreading emerge as fundamental organizing principles that determine what can be simulated, for how long, and with what degree of control [8, 16, 24, 25].

From this perspective, quantum simulation should be understood not as a replacement for classical many-body methods, but as a complementary tool, most naturally suited to non-equilibrium dynamics, short- and intermediate-time evolution, and observables that probe genuinely quantum correlations.

Beyond digital simulation

The approaches discussed in this document represent only one part of a broader landscape. In particular, we have deliberately not addressed in detail analog quantum simulation and experimental quantum emulation, which follow a conceptually distinct route.

In analog quantum simulation, one engineers a physical quantum system whose Hamiltonian closely matches that of the target model. The dynamics of the simulator then directly realizes the dynamics of interest, without decomposition into elementary gates or explicit algorithmic control.

More broadly, many experimental platforms—such as cold atoms in optical lattices, trapped ions, Rydberg arrays, superconducting circuits, or photonic systems—can be used to replicate specific theoretical models and probe their behavior through direct measurement [31–35]. In this paradigm, information about the target system is recovered experimentally rather than computationally.

Simulation versus emulation

From a theoretical standpoint, this experimental strategy can be viewed as an analogue of quantum simulation in a precise sense: instead of encoding the model into a quantum algorithm, one encodes it into the physical properties of a controllable quantum system.

This shift has important consequences:

- the “algorithm” is replaced by Hamiltonian engineering,

- errors arise from imperfect model realization rather than discretization,
- observables are accessed through experimental measurements rather than post-processing.

While this approach sacrifices universality and flexibility, it can offer unparalleled access to specific models and dynamical regimes, often with a degree of realism that is difficult to achieve in digital simulations.

Complementary paradigms

Digital, variational, and analog quantum simulations should not be seen as competing strategies, but as complementary paradigms.

Digital and variational methods emphasize controllability, modularity, and algorithmic transparency. Analog and experimental emulation emphasize physical fidelity and direct access to dynamics.

For many-body physics, the most productive path forward is likely to involve a combination of these approaches, informed by classical simulations and guided by clear physical questions.

Final perspective

Quantum simulation, in all its forms, is best understood as an extension of the many-body toolkit rather than a standalone solution. Its power lies in enabling new ways to explore, approximate, and interpret complex quantum dynamics.

As experimental capabilities and theoretical methods continue to evolve, the distinction between simulation, emulation, and experiment itself may become increasingly blurred. What remains essential is a clear understanding of the assumptions, approximations, and limits inherent in each approach.

With this perspective, quantum simulation becomes not an end in itself, but a means to deepen our understanding of quantum many-body systems.

Bibliography

[1] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010 (cited on page 2).

[2] Riccardo Manenti and Mario Motta. *Quantum Information Science*. Oxford University Press, Aug. 2023 (cited on page 2).

[3] Richard P. Feynman. "Simulating physics with computers". In: *International Journal of Theoretical Physics* 21.6-7 (1982), pp. 467–488. DOI: 10.1007/BF02650179. URL: https://doi.org/10.1007/BF02650179 (cited on pages 5, 37).

[4] Seth Lloyd. "Universal quantum simulators". In: *Science* 273.5278 (1996), pp. 1073–1078. DOI: 10.1126/science.273.5278.1073. URL: https://doi.org/10.1126/science.273.5278.1073 (cited on pages 5, 19, 37).

[5] I. M. Georgescu, S. Ashhab, and Franco Nori. "Quantum simulation". In: *Reviews of Modern Physics* 86.1 (2014), pp. 153–185. DOI: 10.1103/RevModPhys.86.153. URL: https://doi.org/10.1103/RevModPhys.86.153 (cited on pages 5, 37).

[6] Anders W. Sandvik. *Computational Studies of Quantum Spin Systems*. Vol. 1297. AIP, 2010, pp. 135–338. DOI: 10.1063/1.3518900. URL: https://doi.org/10.1063/1.3518900 (cited on pages 6, 9, 10).

[7] Matthias Troyer and Uwe-Jens Wiese. "Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations". In: *Physical Review Letters* 94.17 (2005), p. 170201. DOI: 10.1103/PhysRevLett.94.170201. URL: https://doi.org/10.1103/PhysRevLett.94.170201 (cited on pages 6, 7, 10, 33).

[8] Ulrich Schollwöck. "The density-matrix renormalization group in the age of matrix product states". In: *Annals of Physics* 326.1 (2011), pp. 96–192. DOI: 10.1016/j.aop.2010.09.012. URL: https://doi.org/10.1016/j.aop.2010.09.012 (cited on pages 6, 7, 10, 19, 29, 37).

[9] Steven R. White. "Density matrix formulation for quantum renormalization groups". In: *Physical Review Letters* 69.19 (1992), pp. 2863–2866. DOI: 10.1103/PhysRevLett.69.2863. URL: https://doi.org/10.1103/PhysRevLett.69.2863 (cited on pages 6, 10).

[10] Guifré Vidal. "Efficient simulation of one-dimensional quantum many-body systems". In: *Physical Review Letters* 93.4 (2004), p. 040502. DOI: 10.1103/PhysRevLett.93.040502. URL: https://doi.org/10.1103/PhysRevLett.93.040502 (cited on pages 6, 10).

[11] Pasquale Calabrese and John Cardy. "Evolution of entanglement entropy in one-dimensional systems". In: *Journal of Statistical Mechanics: Theory and Experiment* 2005.04 (2005), P04010. DOI: 10.1088/1742-5468/2005/04/P04010. URL: https://doi.org/10.1088/1742-5468/2005/04/P04010 (cited on pages 6, 7, 10, 19).

[12] Matthew B. Hastings. "An area law for one-dimensional quantum systems". In: *Journal of Statistical Mechanics: Theory and Experiment* 2007.08 (2007), P08024. DOI: 10.1088/1742-5468/2007/08/P08024. URL: https://doi.org/10.1088/1742-5468/2007/08/P08024 (cited on page 10).

[13] Hui Shao and Anders W. Sandvik. "Progress on stochastic analytic continuation of quantum Monte Carlo data". In: *Physics Reports* 1003 (2023). Progress on stochastic analytic continuation of quantum Monte Carlo data, pp. 1–88. DOI: https://doi.org/10.1016/j.physrep.2022.11.002. URL: https://www.sciencedirect.com/science/article/pii/S0370157322003921 (cited on page 10).

[14] Giuseppe Carleo et al. "Localization and Glassy Dynamics Of Many-Body Quantum Systems". In: *Scientific Reports* 2.1 (2012), p. 243. DOI: 10.1038/srep00243. URL: https://doi.org/10.1038/srep00243 (cited on page 10).

[15] Federico Becca and Sandro Sorella. *Quantum Monte Carlo Approaches for Correlated Systems*. Cambridge University Press, 2017 (cited on page 10).

[16] Elliott H. Lieb and Derek W. Robinson. "The finite group velocity of quantum spin systems". In: *Communications in Mathematical Physics* 28.3 (1972), pp. 251–257. DOI: 10.1007/BF01645779. URL: https://doi.org/10.1007/BF01645779 (cited on pages 16, 30, 37).

[17] Matthew B. Hastings and Tohru Koma. "Spectral gap and exponential decay of correlations". In: *Communications in Mathematical Physics* 265.3 (2006), pp. 781–804. DOI: 10.1007/s00220-006-0030-4. URL: https://doi.org/10.1007/s00220-006-0030-4 (cited on pages 16, 30).

[18] Pascual Jordan and Eugene Wigner. ""Uber das Paulische Äquivalenzverbot". In: *Zeitschrift für Physik* 47.9-10 (1928), pp. 631–651. DOI: 10.1007/BF01331938. URL: https://doi.org/10.1007/BF01331938 (cited on page 18).

[19] Sergey Bravyi and Alexei Kitaev. "Fermionic quantum computation". In: *Annals of Physics* 298.1 (2002), pp. 210–226. DOI: 10.1006/aphy.2002.6254. URL: https://doi.org/10.1006/aphy.2002.6254 (cited on page 18).

[20] Dominic W. Berry et al. "Efficient quantum algorithms for simulating sparse Hamiltonians". In: *Communications in Mathematical Physics* 270.2 (2007), pp. 359–371. DOI: 10.1007/s00220-006-0150-x. URL: https://doi.org/10.1007/s00220-006-0150-x (cited on page 19).

[21] Hale F. Trotter. "On the product of semi-groups of operators". In: *Proceedings of the American Mathematical Society* 10.4 (1959), pp. 545–551 (cited on pages 23, 37).

[22] Masuo Suzuki. "Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems". In: *Communications in Mathematical Physics* 51.2 (1976), pp. 183–190 (cited on pages 24, 37).

[23] Marin Bukov, Luca D'Alessio, and Anatoli Polkovnikov. "Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to Floquet engineering". In: *Advances in Physics* 64.2 (2015), pp. 139–226. DOI: 10.1080/00018732.2015.1055918. URL: https://doi.org/10.1080/00018732.2015.1055918 (cited on page 26).

[24] A. I. Larkin and Yu. N. Ovchinnikov. "Quasiclassical method in the theory of superconductivity". In: *Soviet Physics JETP* 28 (1969), pp. 1200–1205 (cited on pages 30, 31, 33, 37).

[25] Juan Maldacena, Stephen H. Shenker, and Douglas Stanford. "A bound on chaos". In: *Journal of High Energy Physics* 2016.8 (2016), p. 106. DOI: 10.1007/JHEP08(2016)106. URL: https://doi.org/10.1007/JHEP08(2016)106 (cited on pages 30, 31, 33, 37).

[26] John Preskill. "Quantum computing in the NISQ era and beyond". In: *Quantum* 2 (2018), p. 79. DOI: 10.22331/q-2018-08-06-79. URL: https://doi.org/10.22331/q-2018-08-06-79 (cited on page 34).

[27] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. "Error mitigation for short-depth quantum circuits". In: *Physical Review Letters* 119.18 (2017), p. 180509. DOI: 10.1103/PhysRevLett.119.180509. URL: https://doi.org/10.1103/PhysRevLett.119.180509 (cited on page 35).

[28] Suguru Endo et al. "Hybrid quantum-classical algorithms and quantum error mitigation". In: *Journal of the Physical Society of Japan* 90.3 (2021), p. 032001. DOI: 10.7566/JPSJ.90.032001. URL: https://doi.org/10.7566/JPSJ.90.032001 (cited on page 35).

[29] Jun Li and Simon C. Benjamin. "Efficient variational quantum simulator incorporating active error minimization". In: *Physical Review X* 7.2 (2017), p. 021050. DOI: 10.1103/PhysRevX.7.021050. URL: https://doi.org/10.1103/PhysRevX.7.021050 (cited on page 37).

[30] Xin Yuan et al. "Theory of variational quantum simulation". In: *Quantum* 3 (2019), p. 191. DOI: 10.22331/q-2019-10-07-191. URL: https://doi.org/10.22331/q-2019-10-07-191 (cited on page 37).

[31] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. "Many-body physics with ultracold gases". In: *Reviews of Modern Physics* 80.3 (2008), pp. 885–964. DOI: 10.1103/RevModPhys.80.885. URL: https://doi.org/10.1103/RevModPhys.80.885 (cited on page 37).

[32] Rainer Blatt and Christian F. Roos. "Quantum simulations with trapped ions". In: *Nature Physics* 8.4 (2012), pp. 277–284. DOI: 10.1038/nphys2252. URL: https://doi.org/10.1038/nphys2252 (cited on page 37).

[33] Antoine Browaeys and Thierry Lahaye. "Many-body physics with individually controlled Rydberg atoms". In: *Nature Physics* 16.2 (2020), pp. 132–142. DOI: 10.1038/s41567-019-0733-z. URL: https://doi.org/10.1038/s41567-019-0733-z (cited on page 37).

[34] Andrew A. Houck, Hakan E. T"ureci, and Jens Koch. "On-chip quantum simulation with superconducting circuits". In: *Nature Physics* 8.4 (2012), pp. 292–299. DOI: 10.1038/nphys2251. URL: https://doi.org/10.1038/nphys2251 (cited on page 37).

[35] Alán Aspuru-Guzik and Philip Walther. "Photonic quantum simulators". In: *Nature Physics* 8.4 (2012), pp. 285–291. DOI: 10.1038/nphys2253. URL: https://doi.org/10.1038/nphys2253 (cited on page 37).