



Where Quantum and Fluids entangle

Quantum Computational Fluid Dynamics Scaling and Memory Requirements

Project number 101080085

Call:	HORIZON-CL4-2021-DIGITAL-EMERGING-02
Topic:	HORIZON-CL4-2021-DIGITAL-EMERGING-02-10
Type of action:	HORIZON Research and Innovation Actions
Granting authority:	European Commission–EU
Project starting date: fixed date:	1 November 2022
Project end date:	31 October 2026
EU–Project duration:	48 months
Project Coordinator:	University of Hamburg (UHH)
Work Package Leader:	University of Hamburg (UHH)
Cooperations:	Technical University of Crete (TUC) PlanQC
Deliverable number:	D4.2
WP contributing to the deliverable:	WP#4 Tensor Network Simulations
Deliverable Type:	Report
Revision:	2
Dissemination level:	Public
Due Submission date:	21.04.2026
Prepared By:	UHH, TUC, PlanQC
Internal Reviewers:	DJ, DA, MK
Final Approval:	DJ



This project receives funding from the European Union's Horizon 2020 / Horizon Europe research and innovation actions programme under Grant Agreement #101080085.

Revision History

Version	Date	Who	Changes
R0	10.10.2025	NH	
R1	13.04.2026	NH	<p>Rayleigh–Bénard Convection:</p> <ul style="list-style-type: none"> • Grid–Rayleigh-number selection for the RB test cases has been clarified, referencing the boundary layer resolution requirements of [1]. <p>S-Bend:</p> <ul style="list-style-type: none"> • The grid-Re coupling for the S-bend test case has been clarified, and the grid-growth exponent β has been corrected. • The asymptotic advantage conditions in Section 2.1.1 have been updated to reflect the corrected β, and exceptions at $\varepsilon_F = 10^{-10}$ are now stated. • Table 2.14 has been updated accordingly.
R2	21.04.2026	NH	<ul style="list-style-type: none"> • Citation [2] has been added to the appropriate places in Section 1.1 and 3.1. • Citation of our recent preprint [3] on Rayleigh-Bénard convection has been added to the appropriate Section 2.1.3.



Executive summary

This deliverable establishes scaling requirements for variational quantum algorithms (VQAs) for computational fluid dynamics (CFD) by using tensor networks as a classical proxy to derive empirical lower bounds on the quantum resources needed at a given accuracy. We assess the representability of CFD solutions in matrix product state (MPS) format by compressing fields produced by our in-house finite-difference code and the industry-standard OpenFOAM solver, Helyx. To enable fair, quantitative comparison, we introduce an advantage factor A that combines the grid-resolution requirement with a model for how the required MPS bond dimension χ —the parameter that sets the computational and storage requirements—scales with Reynolds number Re . Using three case studies—laminar cylinder flow, Rayleigh–Bénard convection (transition to turbulence), and an engineering S-bend, we map how χ grows with Re to achieve a fixed accuracy and translate this into projected time and memory costs. The results indicate that, wherever the physics admits compact tensor representations (moderate χ), MPS offers clear reductions in memory footprint and runtime at fixed error compared with classical baselines. From a quantum-algorithm perspective, the tensor-network study is a scaling analysis that yields classical lower bounds on the achievable speedup by the quantum algorithm. We do not expect those asymptotics to change in quantum implementations. What matters for near-term viability are the constant prefactors, since quantum computing resources are scarce. We therefore focus on optimizing the quantum algorithms’ efficiency: prototyping VQAs on 1D benchmarks (nonlinear Schrödinger equations and Burgers’ equation), using MPO-to-circuit compilation and the QNPU concept for nonlinear terms, and exploring depth reduction (e.g., trading depth for qubits with mid-circuit measurement and classical control), ansatz and optimizer choices, and barren-plateau mitigation. These studies target smaller prefactors through shallower circuits, fewer optimization parameters, and more stable training. All data supporting this report are publicly available; newly generated datasets (including figure data) are archived under DOI: 10.25592/uhhfdm.18026 in the UHH research data repository.



Contents

1.1	Introduction	6
1.2	Fundamentals of Data Description	7
1.2.1	Accuracy	7
1.3	Assessment Criteria	7
2.1	Tensor Networks	8
2.1.1	Theoretical Scaling Analysis (2D)	9
2.1.2	Case Study I: Flow cylinder using curvilinear coordinates	11
2.1.3	Case Study II: Rayleigh-Bénard Convection	13
2.1.4	Case Study III: Flow in an S-Bend	16
2.1.5	Discussion	17
3.1	Variational Quantum Algorithms	18
3.1.1	Amplitude Encoding and Quantum Nyquist-Shannon Theorem	18
3.1.2	VQA vs TN: Scaling Analysis	20
3.1.3	Tensor-Programmable VQA	21
3.1.4	Case Study IV: Nonlinear Schrödinger equation	22
3.1.5	Case Study V: Burgers' equation	23
3.1.6	Case Study VI: Circuit Optimization	25
3.1.7	Quantum Sparse Optimization Near the Edge of Chaos	26
3.1.8	Discussion	28
4.1	Conclusion	29



List of abbreviations

Acronym / Short Name	Meaning
BC	Boundary Condition
BEC	Bose-Einstein Condensates
CA	Consortium Agreement
CFD	Computational Fluid Dynamics
D	Deliverable
DMP	Data Management Plan
DNS	Direct Numerical Simulation
DOI	Digital Object Identifier
EC	European Commission
ENG	ENGYS SRL
EU	European Union
FD	Finite Difference
FDR	Research Data Repository
FV	Finite Volume
FZL	Jülich Research Centre
KPI	Key Performance Indicator
MPO	Matrix Product Operator
MPS	Matrix Product State
NLSE	Nonlinear Schrödinger Equation
NVPS	Number of Variables Parameterizing the Solution
PlanQC	PlanQC GmbH
PO	Project Officer
QCFD	Quantum Computational Fluid Dynamics
QNPU	Quantum Nonlinear Processing Unit
RMSE	Root Mean Squared Error
ROM	Reduced Order Models
SVD	Singular Value Decomposition
TCI	Tensor Cross Interpolation
TEE	Topological Entanglement Entropy
TN	Tensor Network
TT	Tensor Train
TUC	Technical University of Crete
TUHH	Technical University of Hamburg
TUM	Technical University of Munich
UHH	University of Hamburg
VQA	Variational Quantum Algorithm
WP	Work Package
WPL	Work Package Leader

1.1 Introduction

Resolving the large ranges of spatial and temporal scales in computational fluid dynamics (CFD) poses a significant challenge for classical hardware due to intensive computational and energy demands. The advent of quantum computing presents a new paradigm for tackling these problems, but a rigorous analysis of its resource scaling is essential to assess its potential. This report analyzes and benchmarks the scaling and memory requirements of tensor network (TN) methods, specifically matrix product states (MPS), and variational quantum algorithms (VQAs) on representative CFD problems. The focus is on quantifying these metrics to establish the achievable lower bound on quantum speedup.

Building on the MPS algorithms developed within WP4, we benchmark their performance against established classical solvers: our in-house finite-difference (FD) code and the industry-standard OpenFOAM solver, Helyx [4]. To provide a quantitative basis for comparison, we introduce an advantage factor (A) in Section 2.1.1. It is determined by modeling the required bond dimension (χ) as a power-law function of the Reynolds number (Re), which quantifies the flow's complexity. By integrating this scaling with the required grid resolution, we generate clear performance metrics for both time and memory efficiency. We apply this analysis to the following three representative use-cases:

1. Laminar flow past a cylinder (Section 2.1.2).
2. Rayleigh–Bénard convection (transition to turbulence) (Section 2.1.3) .
3. Flow through an S-bend (a realistic engineering geometry) (Section 2.1.4).

These benchmarks are selected in order to cover a progression from laminar to turbulent dynamics and to include a practical industrial configuration.

The first two cases are well documented in the literature [5, 6], enabling rigorous comparison and validation and the last one was specified in the initial project proposal. This strategy provides a robust scientific baseline for validation before extending our methods to more complex engineering geometries. For each case, we generated steady-state solutions across a range of Reynolds numbers (Re) and grid sizes (N) to determine the bond dimension (χ) required to represent the solution at a fixed accuracy. Establishing how χ scales with Re is crucial, as the computational cost of the MPS algorithms is governed by χ as detailed in the upcoming section. Based on this analysis, we demonstrate that an MPS-based approach offers a promising path to substantial memory and runtime reductions at a fixed accuracy compared to classical algorithms. Further discussion is provided in Section 2.1.5.

In the QCFD framework, TNs serve as the classical foundation for building and understanding variational quantum circuits. MPS and matrix product operators (MPO) provide compact, variational representations of flow fields and operators that can be directly mapped to shallow quantum circuits implementing the same functional structure. This correspondence allows TN simulations to act as high-level, classical analogues of the VQAs, sharing the same optimization principles while offering a tractable way to test scaling, expressivity, and accuracy. Within this framework, the Quantum Nonlinear Processing Unit (QNPU) is a quantum circuit that extends the standard VQA architecture by enabling the evaluation of nonlinear cost functions directly on quantum hardware: it combines multiple copies of variational states and applies problem-specific linear operators to realise nonlinear terms that appear, for example, in the Navier–Stokes equations. In practice, the operators entering the QNPU are first expressed as MPOs and then translated into gate-level quantum circuits, establishing a direct pathway from classical TN formulations to variational quantum implementations. Consequently, TN simulations not only benchmark and validate the quantum approach but also provide lower bounds on the achievable quantum advantage [7], with the quantum networks expected to reproduce the same variational behavior at significantly reduced computational cost.

We begin in Section 3.1.1 by introducing the fundamental concepts of VQAs alongside a quantum version of the Nyquist-Shannon theorem [2], which provides a theoretical bound on the circuit width required to represent a function state. In Section 3.1.2, we demonstrate the potential of VQAs in representing 3D turbulent fields where we show that VQAs require fewer parameters than the MPS representation. Furthermore, through scaling arguments specific to this case, we expect VQAs to outperform the MPS-based approach in the limit of large system size N , yielding a significant memory as well as computational gain. In Section 3.1.3, we discuss the general route from the MPS format to quantum algorithms by compiling differential operators expressed as MPOs into quantum circuits. In Sections 3.1.4 and 3.1.5, we benchmark VQAs on typical one-dimensional problems, namely:

1. Nonlinear Schrödinger equation. (Section 3.1.4)



2. Burgers' equation. (Section 3.1.5)

While we do not anticipate a change in the fundamental scaling laws, our current analysis does not account for constant prefactors. These prefactors are critical for practical applications and can be substantially improved through algorithmic optimization, encompassing everything from the ansatz design to the classical optimizer. The purpose of these initial, small-scale studies is two-fold: first, to benchmark the necessary quantum hardware resources, and second, to identify and mitigate configurations that may lead to barren plateaus. In short, the simple case studies are designed to refine the VQA design and gather crucial engineering data before attempting large, potentially costly, simulations. Quantum circuits tend to become deeper with increasing complexity of the task, making them difficult to execute on current noisy hardware. Reducing circuit depth is therefore a key design goal. In Section 3.1.6, we discuss a strategy that trades circuit depth for qubit count, namely by using additional qubits together with mid-circuit measurements and classical control to emulate the same transformation with a shallower, effectively non-unitary circuit. Finally, in Section 3.1.7, we propose an approach to mitigate the challenges posed by barren plateaus. This severely limits the trainability of existing VQAs, leading to exponentially vanishing gradients that drastically inflate the prefactors governing the computational cost of the optimization process.

To ensure full reproducibility, all data presented in this report is publicly available. Data drawn from previous work is cited at its original source. The new datasets generated specifically for this study, including all data used to create the figures, are archived and can be accessed via the following Digital Object Identifier (DOI): 10.25592/uh-hfdm.18026 according to the submitted Data Management strategy, in a dedicated research data repository (FDR) hosted at the University of Hamburg (UHH), <https://www.fdr.uni-hamburg.de/communities/qcfd/>, and associated within the project community QCFD.

1.2 Fundamentals of Data Description

This section's scope is to secure data quality. The overall goal of QCFD is to combine the computational resources of quantum hardware with CFD industrial applications. To this end, established CFD tools, that have proven to show a good agreement with the underlying physical phenomena, are used as benchmark (verification) examples for the algorithmic transition from classical to the quantum frameworks. To highlight the application relevance, the document distinguishes between scientific and industrial motivation. Each example case holds a simplified categorization which transits between industrial and scientific relevance, cf. Figure 1.1.




Figure 1.1: Application indicator slider

1.2.1 Accuracy

Data accuracy is fundamental for expressive and trustful benchmarking and is therefore also a key aspect of this deliverable 4.2. In this regard, the employed guidelines are given next:

- List of file formats with the corresponding accuracy, accessibility without compression
 - .csv
 - .npy
 - Double-precision floating-point number format
 - Each component of a complex number retains full double-precision accuracy (53-bit significand, \approx 16 decimal digits).

1.3 Assessment Criteria

The assessment criteria are defined as follows:



- **L_2 norm and relative L_2 error.** Let $\Omega \subset \mathbb{R}^d$ be the spatial domain and $\mathbf{x} \in \Omega$ the spatial coordinate. For a field $u : \Omega \rightarrow \mathbb{C}$ (or \mathbb{R}), define

$$\|u\|_{L_2(\Omega)} = \left(\int_{\Omega} |u(\mathbf{x})|^2 d\Omega \right)^{1/2}. \quad (1.1)$$

Given $u_1, u_2 : \Omega \rightarrow \mathbb{C}$ (or \mathbb{R}), the relative L_2 error of u_1 with respect to the reference u_2 is

$$\epsilon_{L_2} = \frac{\|u_1 - u_2\|_{L_2(\Omega)}}{\|u_2\|_{L_2(\Omega)}}. \quad (1.2)$$

- **Fidelity (F):** Often understood as a similarity measure between two quantum state vectors. In this context, we adopt the conventional measure of squared fidelity (F^2) used in TN simulations and quantum information theory. It is calculated by taking the squared absolute value of the scalar product between two fields, u_1 and u_2 , normalized by the product of their squared norms. It reads as

$$F = \frac{|u_1^\dagger u_2|^2}{\|u_1\|^2 \cdot \|u_2\|^2}, \quad (1.3)$$

where $(\cdot)^\dagger$ denotes the conjugate transpose.

We further define the infidelity, the deviation from a fidelity of 1, by $\epsilon_F = 1 - F$.

- **The Root Mean Squared Error (RMSE):** Denoted by ϵ_{RMSE} , this metric quantifies the quadratic mean of the differences between two fields, u_1 and u_2 . It represents a standard measure of absolute error, where the differences are squared before they are averaged. The RMSE is defined as

$$\epsilon_{\text{RMSE}} = \sqrt{\frac{1}{N} \sum_{i=1}^N (u_{1,i} - u_{2,i})^2} \quad (1.4)$$

2.1 Tensor Networks

In the following, we first discuss the TN formalism, specifically the MPS which is inherently compatible with quantum algorithms. An MPS is a factorization of an n -index tensor into a chain of three-index tensors connected by virtual indices whose dimension is the bond dimension (see panel 3 of Fig. 2.2). We aim to use this formalism to formulate and simulate quantum-inspired algorithms for CFD, and thereby assess in which use cases and regimes such methods may offer an advantage compared to classical methods. Since an MPS is a parametrized representation of a state, such as a fluid field, its expressivity can be limited by the number of allowed parameters. This is illustrated in the third panel of Fig. 2.2. The bond dimension χ at the shared indices between MPS tensors precisely controls this number of parameters (i.e., tensor entries), as well as determines how an MPS-based algorithm scales, both in computational time and in required memory. It is also evident that this control parameter governs the expressivity of the MPS: more complex flow scenarios typically will require greater expressivity (a larger bond dimension).

To assess whether an MPS algorithm can have an advantage in terms of runtime T and storage M compared to classical algorithms, we must estimate how the bond dimension grows with system size and with flow complexity. An advantage arises if this growth, and the associated algorithmic complexity, is slower than for a conventional algorithm. To quantify this, in Section 2.1.1 we show how an MPS algorithm scales with bond dimension and relate system size to flow complexity (e.g., via the Kolmogorov length scale) using the Reynolds number Re as a driver. Moreover, we establish a relationship between bond dimension and flow complexity that can be evaluated with empirical data from classical simulations. This yields a ratio of expected classical runtime and memory to the corresponding MPS values.

We then analyze this ratio across three use cases: (I) flow around a cylinder in curvilinear coordinates (Section 2.1.2); (II) Rayleigh-Bénard convection (Section 2.1.3); and (III) an S-bend (Section 2.1.4). These examples were chosen to span different problem classes: a study in curvilinear coordinates with potential industrial relevance, an academic case probing turbulence structure, and a clearly industrial geometry (the S-bend). The key conclusion is that scaling statements for the MPS are use-case dependent; moreover, there may be regimes in which



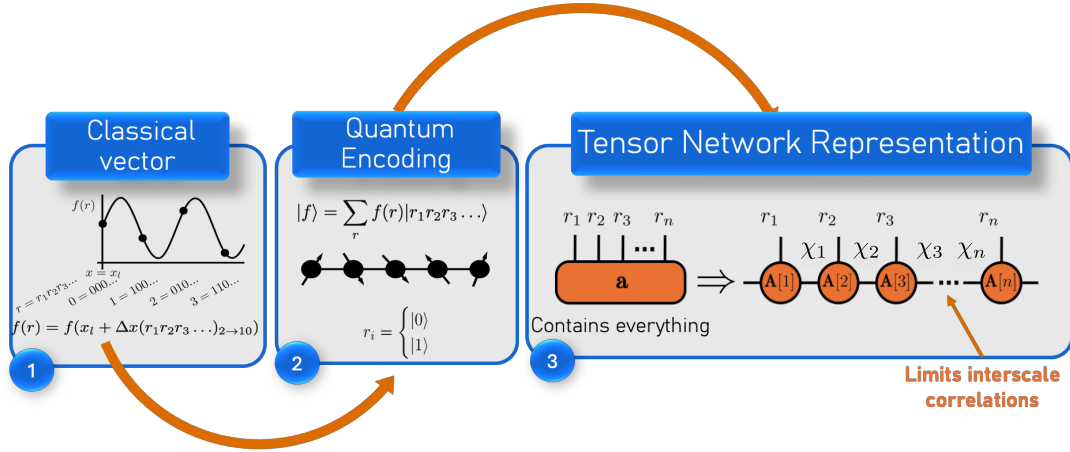


Figure 2.2: (1) Represents the classical vector, containing the discretized field values the conventional solver operates upon. (2) Through Quantum encoding, this classical vector can be expressed into a compressed form that can be ported to a quantum computer. (3) In this format, it can be thought of as decomposing a high-dimensional object (big tensor \mathbf{a}) into a chain of smaller tensors, connected through an index χ , whose size dictates the achievable compression and efficiency.

the bond dimension saturates at higher complexity (as reported by [8]), and where scaling based on mathematical fidelity measures does not fully reflect the resources required to resolve the relevant physics.

Operation	Method	Leading cost (after rounding to target rank)
Addition $c = a + b$	Direct	$O(\chi^3)$
	Variational	$O(\chi^3)$
Hadamard $c = a \odot b$	Direct	$O(\chi^6)$
	Variational	$O(\chi^4)$
	Tensor-Train Multiplication [9] (TTM)	$O(\chi^3)$
Matvec $b = \underline{\mathbf{M}}a$	Direct	$O(\chi_o^3 \chi^3)$
	Variational	$O(\chi_o \chi^3)$
Linear solve $\underline{\mathbf{M}}a = b$	Variational with local exact LU	$O(\chi^6)$
	Variational with local iterative (fast matvec)	$O(\chi_o \chi^3 + \chi_o^2 \chi^2)$

Table 2.1: Scaling table for core MPS algorithmic operations. Here, χ_o denotes the bond dimension of the MPO $\underline{\mathbf{M}}$, and χ the bond dimension of a , b , and c (assumed comparable for simplicity) [10].

2.1.1 Theoretical Scaling Analysis (2D)

We consider a discretization on a 2D system with N_x and N_y being the number of grid points along the x-axis and y-axis respectively, with total number of points being $N = N_x N_y$. We stress that, in order to encode our field into an MPS, we always choose the number of points along x and y to be strictly $N_x = 2^{n_x}$, $N_y = 2^{n_y}$, where $n_x \geq 1$, $n_y \geq 1$. When encoding the fields into an MPS (see [10] for details), the decomposition of an arbitrary vector \vec{a} , reshaped as 2^n -component tensor $\mathbf{a}_{r_1, r_2, \dots, r_n}$, into the MPS format is done by performing a series of singular value decomposition (SVDs), where $\mathbf{A}[i]$ are order-3 tensors with mode indices $r_i \in \{0, 1\}$, and bond indices χ_i . Fig. 2.2 summarises the steps required to transform a vector field \vec{a} onto an MPS. Alternatively Tensor Cross interpolations (TCI) methods can be utilized [11]. In MPS formalism, the bond dimension χ controls how much spatial and temporal structure/complexity can be represented. As the Reynold number Re increases, the structure and complexity of the field is expected to increase. In the next part, we lay out the mathematical framework used to assess the suitability of MPS-based solvers in comparison to classical solvers, focusing on the scaling exponent and ignoring constant prefactors.



The computational resources for the simulation are dictated by the size of the MPS representation, which is expected to grow with the flow complexity, i.e. as Re increases. We quantify this relationship by modeling the bond dimension χ and the necessary grid size N as functions of Re . Following Appendix E of [7, 8], we adopt the scaling law $\chi \sim \text{Re}^\alpha$, and further $N \sim \text{Re}^\beta$.

Here, α is estimated from required MPS bond dimensions χ at fixed accuracy for different system sizes, and β reflects the necessary grid size to resolve all involved length scales for a specific Re . While the specifics of the relationship between the Re and χ are problem-dependent, a larger χ is typically required to capture the intricate, multi-scale fluid correlations as Re increases. This link is central to our analysis because the computational cost for the MPS scales with χ , and is dominated by $O(\chi^3)$ scalings (see Table 2.1 for details). Since cubic scaling is characteristic of several key algorithmic steps, including matrix-vector (Matvec) multiplication and Hadamard product, we select Matvec as a representative benchmark for modeling overall algorithm performance of the MPS.

Applying these scaling assumptions to our benchmark allows us to compare the MPO-MPS approach with a classical FD method. A sparse FD matrix-vector multiplication scales with the number of grid points N . Given that $N \sim \text{Re}^\beta$, the complexity can be expressed in terms of the Reynolds number as Re^β , whereas an MPO-MPS contraction scales as $\chi_o \text{Re}^{3\alpha} \log \text{Re}$, where χ_o denotes the MPO bond dimension (for standard FD stencils on uniform grids, $\chi_o = O(1 - 10)$). Memory (M) shows an analogous comparison: $M_{\text{FD}} \sim \text{Re}^\beta$ versus $M_{\text{MPS}} \sim \text{Re}^{2\alpha} \log \text{Re}$ for a single scalar field. A summary is presented in Table 2.2.

Quantity	FD	MPS
Runtime T (Matvec)	$O(N) \sim \text{Re}^\beta$	$O(\chi_o \chi^3) \sim \chi_o \text{Re}^{3\alpha}$
Memory M	$O(N) \sim \text{Re}^\beta$	$O(\chi^2) \sim \text{Re}^{2\alpha}$

Table 2.2: This table summarizes the runtime T and storage cost M in both FD and MPS approach. Note that for Matvec operation in the MPS format, we choose the variational method with leading cost given in Table 2.1 and we neglect constant prefactors.

Whether the algorithm achieves a true speedup depends on which resource, arithmetic computation or data movement (memory), is the bottleneck on the target platform. In many CFD problems, classical solvers are highly memory-bound ($\sim 95\%$ time spent on memory traffic). In this scenario, reducing the memory footprint and, consequently, data movement can produce significant speedups, even if the number of arithmetic operations rises. Conversely, on compute-limited hardware, reducing the arithmetic count dominates performance. Given that the balance between these factors is platform-specific, we avoid a single metric and instead report the runtime advantage factor (A_T) and memory advantage factor (A_M) separately, where A characterizes the ratio of FD to MPS cost. To summarize, we define advantage factors (A) as the ratio of FD to MPS cost (time or memory). Neglecting polylogarithmic terms and constant prefactors, we get

$$A_T \sim \text{Re}^{\beta-3\alpha}, \quad A_M \sim \text{Re}^{\beta-2\alpha}.$$

Thus, two practical decision criteria are identified: (i) if $\alpha < \beta/3$, the MPS formulation is asymptotically favorable in runtime; (ii) if $\alpha < \beta/2$, it is asymptotically favorable in memory. These criteria will be evaluated in following case studies in Section 2.1 by estimating (α, β) from the data and reporting time and memory at matched accuracy.

2.1.2 Case Study I: Flow cylinder using curvilinear coordinates

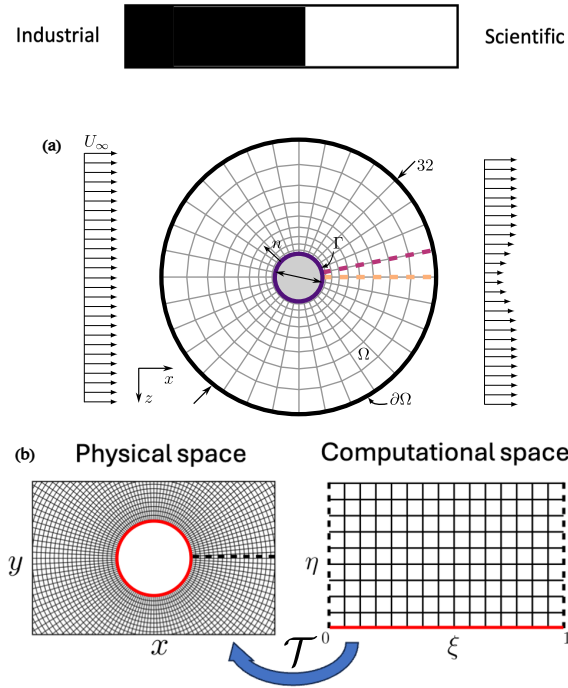


Figure 2.3: (a) A unit-diameter cylinder ($D = 1$) with boundary Γ is centered in a circular domain Ω whose outer boundary $\partial\Omega$ has diameter $32D$ (i.e., $32D$). Grid lines are shown in light gray; the exterior boundary $\partial\Omega$ is indicated in black and the interior cylinder boundary Γ in magenta. (b) Transformation between computational and physical domain via mapping \mathcal{T} .

Flow past a circular cylinder is a canonical CFD benchmark with extensive prior literature [5] and, in certain regimes, available analytical results. Despite its geometric simplicity, it exhibits rich unsteady dynamics such as the von Kármán vortex street at moderate Reynolds numbers. We discretize the domain using a body-fitted curvilinear grid characterized by the smooth mapping: $(\xi, \eta) \in [0, 1]^2 \mapsto (x, y)$, so that the curved physical boundaries are mapped to straight boundaries in the rectangular computational domain, as illustrated in Fig. 2.3b. Further details on the setup can be found in [10]. Here, we aim to assess how well MPS can represent the fluid fields on this curvilinear grid (or likewise on the computational domain). To this end, we run FD simulations to a steady state, encode the final fields into MPS, and study the required bond dimension for each state variable.

Scaling Analysis for Laminar Flows

We use the curvilinear body-fitted mesh described in [10]. The Reynolds number Re for each grid size N is selected based on the requirement for a fixed boundary-layer resolution and summarized in Table 2.3.

N	8,192	32,768	131,072	524,288
Re	1.1	7.9	41.7	191.1

Table 2.3: This table shows the growth of the number of points in our grid with increasing Re .

From this table, we estimate the growth exponent β to be approximately 0.8. For each pair (N, Re) , we perform a steady-state simulation using an in-house FD solver, and determine the minimal MPS bond dimension $\chi(\varepsilon; \text{Re})$ with tolerance ε_{L_2} which is the relative L_2 error. In Fig. 2.4, we look at the scaling obtained when representing u, v and p into an MPS at three different ε_{L_2} , we observe that the scaling depends strongly on the selected ε_{L_2} .

Taking $\beta = 0.8$ and considering the two criteria derived in Section 2.1.1, we obtain $\alpha < \beta/3 = 0.266$ for runtime advantage and $\alpha < \beta/2 = 0.4$ for memory advantage. The results are summarised in Fig. 2.5 for the data in Table 2.4.

In this particular example, we show that an MPS algorithm already provide a memory advantage in this simple flow regime. Concerning the runtime T for the Matvec operation, the MPS algorithm shows an advantage for

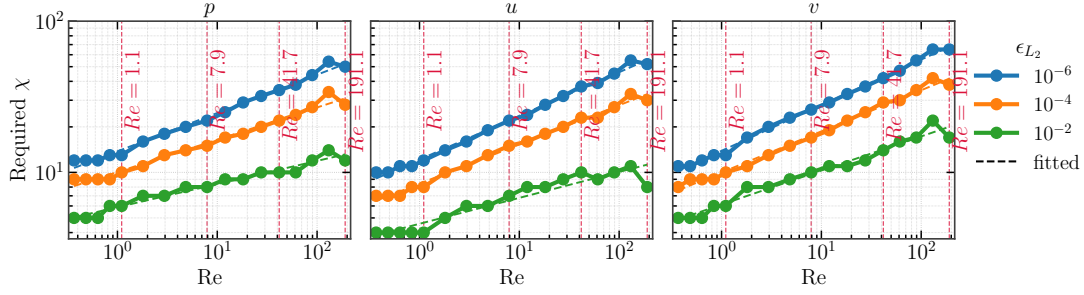


Figure 2.4: Required MPS bond dimension χ as a function of Re to achieve a prescribed relative L_2 error ϵ_{L_2} for each field (p, u, v). Solid curves show the minimal χ obtained from the data (log–log axes); dashed curves are least-squares power-law fits $\chi \approx C Re^\alpha$ for each ϵ_{L_2} and field. Vertical red dashed lines indicate reference Re values, highlighting the grid transition.

Field	$\epsilon_{L_2} = 10^{-6}$	$\epsilon_{L_2} = 10^{-4}$	$\epsilon_{L_2} = 10^{-2}$
p	$\alpha \approx 0.256$	$\alpha \approx 0.219$	$\alpha \approx 0.159$
u	$\alpha \approx 0.292$	$\alpha \approx 0.265$	$\alpha \approx 0.192$
v	$\alpha \approx 0.304$	$\alpha \approx 0.277$	$\alpha \approx 0.235$

Table 2.4: Scaling fits $\chi \approx Re^\alpha$ at fixed ϵ_{L_2} .

Field	$\epsilon_{L_2} = 10^{-6}$	$\epsilon_{L_2} = 10^{-4}$	$\epsilon_{L_2} = 10^{-2}$
p	Runtime	Runtime	Runtime
	Memory	Memory	Memory
u	Runtime	Runtime	Runtime
	Memory	Memory	Memory
v	Runtime	Runtime	Runtime
	Memory	Memory	Memory

Figure 2.5: Comparison of the computed exponent α : green color indicates that the MPS representation has an advantage, while red color indicates that the MPS representation does not show a benefit.

$\epsilon_{L_2} = 10^{-2}$, however, achieving an advantage becomes more challenging when aiming for a smaller relative L_2 error $< 10^{-4}$. It is important to stress that, as shown in Fig. 2.4, the bond dimension is still growing, however, from [8] where work has been done on turbulent DJ/DT flows, they observed that in their particular example, the bond dimension saturated after $Re > 10^4$. Further studies are necessary to determine whether this behavior also applies to our case. Nevertheless, it is remarkable that MPS already show some promise in this simple example and in Section 2.1.5, we will continue the discussion on the advantage of memory-efficient algorithm.

Compressibility Analysis of the MPOs representing Curvilinear Operators

It is important to emphasize that the overall complexity also depends on the bond dimensions of the MPOs representing the differential operators. Accordingly, Fig. 2.6a. illustrates the bond dimension of the curvilinear operators as MPOs (dashed line) compared to the corresponding operators defined on a Cartesian mesh (solid line). As it is evident in this figure, the generalized operators require a comparable bond dimension, increasing only slightly for larger grid sizes. Furthermore, the number of NVPS of the MPS representation ($NVPS_{MPS}$) exhibits a remarkable reduction – up to three orders of magnitude – compared to the conventional sparse matrix representation (cf. Fig. 2.6 (b)). For this comparison, we consider a sparse matrix size of $NVPS_{FD} = 5N$ that represents the central point and its four neighboring points in the second-order central FD approximation of the Laplacian operator. Diagonal neighbors are not considered, due to the orthogonality of the grid.

Timing for Transients Flows

Fig. 2.7 provides a direct comparison between the wall-clock time of our MPS-based approach and the classical FD solver. For the system sizes analyzed, both methods exhibit similar computational times. As the grid size increases, the computational time for the MPS simulations shows a constant linear growth rate, only resulting



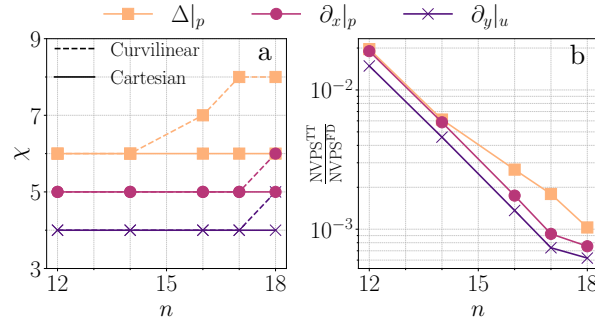


Figure 2.6: Compressibility analysis of the MPOs representing curvilinear operators. Panel (a) depicts the required bond dimension of the curvilinear and Cartesian operators. Panel (b) shows the ratio of number of variables parameterizing the Solution (NVPS) needed to represent the curvilinear operators in MPO format and sparse format. Please note that the lines connecting the data points are added solely for better visualization of the emerging trends. Plots taken from [10].




Figure 2.7: Scaling of the average wall-clock time per iteration (sampled every 200th iteration) for the $Re = 150$ simulations at different grid sizes discussed in [10].

in a simple offset along the y -axis with increasing bond dimension χ . In contrast, the theoretical computational complexity of the classical FD method scales exponentially with the number of qubits. Despite differences in the equation-solving strategies applied to the two methods (e.g., lack of preconditioning in the FD solver), these conclusions remain robust. Since industrial CFD problems routinely involve systems equivalent to over 24 tensors (approximately 20 million degrees of freedom), we project that the MPS-based tensor network method will deliver substantial computational gains, with the break-even point shifting further in favor of the MPS method as further classical solver optimizations are introduced in the MPS methods [10].

2.1.3 Case Study II: Rayleigh-Bénard Convection



In this section, we closely follow the study presented in [12]; we use the same problem setup, discretization, and evaluation protocol. The corresponding results have been submitted for publication and are available as a preprint [3]. Rayleigh-Bénard convection is a classic canonical system for investigating pattern formation, self-organization, and the onset of turbulence in fluids [13, 14]. It occurs when a fluid layer is heated from below and cooled from above, where the resulting vertical temperature gradient drives coherent convective motion [15]. Temperature contrasts sustain many turbulent flows in nature and technology—including stellar and planetary interiors [16], atmospheric circulation [17], and ocean dynamics [18]. From an engineering standpoint, convection is central to heat exchangers and to passive cooling of electronic devices and computer chips.

As illustrated in Fig. 2.8, the canonical Rayleigh-Bénard setup consists of a horizontal fluid layer of thickness D in a gravitational field g , bounded above and below by rigid, thermally conducting plates. Analyses typically assume an infinite (or very large) lateral extent W relative to the vertical dimension [14]. When the imposed




Figure 2.8: Schematic representation of a two-dimensional Rayleigh-Bénard system, illustrating Bénard cells as oval shaped structures with arrows indicating the fluid velocity. This figure was taken from [12].

temperature difference $\Delta\theta$ across the plates exceeds a critical value, the hydrostatic equilibrium—i.e., the hydrostatic state becomes unstable, and organized convection appears—typically as rolls or hexagonal cells [19]. The Rayleigh number Ra quantifies the competition between buoyant driving and viscous/thermal diffusion [20]; it is the principal control parameter governing both the onset and the vigor of convection and is given by

$$Ra = \frac{\beta g D^3 \Delta}{\kappa \nu}, \quad (2.5)$$

where Δ , β , κ and ν are the temperature difference between the two plates, thermal expansion coefficient at constant pressure, thermal diffusivity and momentum diffusivity respectively.

Grid description and Turbulent Temperature Field

The placement of variables on the grid strongly affects stability and accuracy. Collocated layouts (storing pressure and velocity at the same nodes) can cause pressure–velocity decoupling, yielding the “checkerboard” pressure oscillations [5]. To avoid this, we use a staggered grid following Harlow and Welch [21].

Fig. 2.9 compares temperature fields across parameter regimes, highlighting how the Rayleigh number Ra shapes the convective structures. At $Ra = 10^8$, the flow is markedly more turbulent yet still organized by a

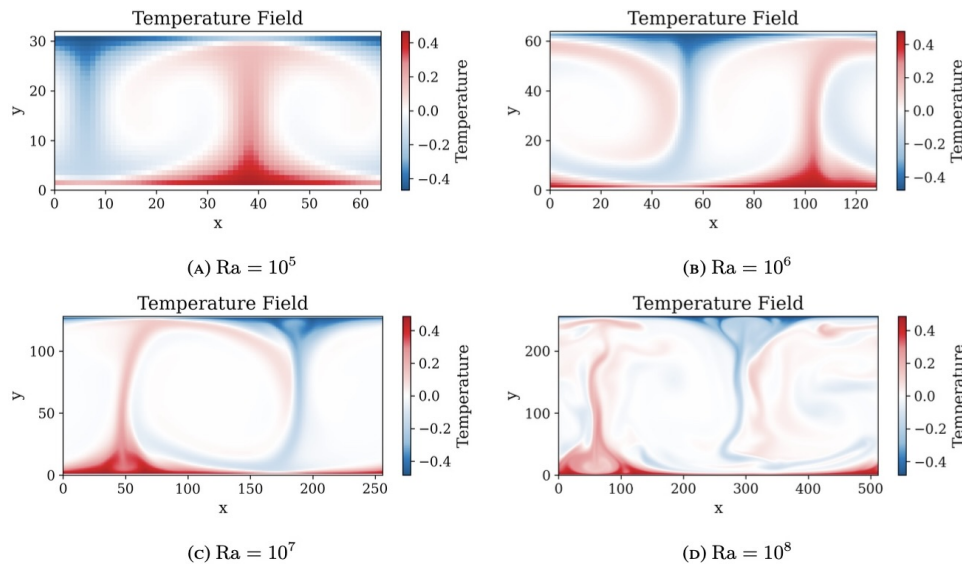


Figure 2.9: Temperature fields for increasing Rayleigh numbers. These plots were taken from [12].

coherent large-scale cellular pattern. The large cells are filled with chaotic substructure and stronger small-scale mixing. Boundary layers are extremely thin, and the dynamics are strongly time-dependent. Multiple length scales coexist, from domain-scale circulations to fine turbulent features. Therefore, in what follows we focus on the temperature field θ , as it is the most challenging field to compress with an MPS.

$N = N_x \times N_y$	64×32 (2048)	128×64 (8192)	256×128 (32768)	512×256 (131072)
Re	316	1000	3162	10000

Table 2.5: Summary of grid size N with Reynolds number Re

Scaling analysis for Laminar-Turbulent Flows

We investigate Rayleigh–Bénard convection in the laminar–turbulent transition and evaluate it with the scaling framework introduced in Section 2.1.1. The mapping from grid size to Reynolds number Re used in our Rayleigh–Bénard tests is listed in Table 2.5 adopted from the supplemental material of [1], where the grid is chosen to resolve thermal and kinetic boundary layers. There, grid–Rayleigh-number pairs are provided for Prandtl number $Pr = 1$, which is also used in the present study. Since $Ra = Re^2$ at $Pr = 1$, the physically motivated grid scaling in terms of Re is directly inherited from the boundary layer resolution requirements of [1]. From these pairs we read out the grid–growth exponent as $\beta \approx 1.2$. For each (Re, N) in Table 2.5, we run a classical FD DNS to a (statistically) steady state, encode the final temperature field θ as an MPS, and determine the required bond dimension χ required to meet a prescribed relative L_2 error ε_{L_2} . The resulting χ -versus-Re curves are shown in Fig. 2.10, and the corresponding fitted exponents α in the relation $\chi \sim Re^\alpha$ are summarized in Table 2.6.

Based on the criteria of Section 2.1.1, an asymptotic runtime advantage requires $\alpha < \beta/3 = 0.4$, which is not satisfied by the fitted values in Table 2.6. An asymptotic memory advantage requires $\alpha < \beta/2 = 0.6$. This is satisfied for $\varepsilon_{L_2} = 10^{-4}$ and $\varepsilon_{L_2} = 10^{-2}$, but α just meets the minimum threshold at $\varepsilon_{L_2} = 10^{-6}$.




Figure 2.10: Comparison of the required bond dimensions to achieve a given ε_{L_2} in the field θ .

ε_{L_2}	1.00×10^{-6}	1.00×10^{-4}	1.00×10^{-2}
α	0.60	0.55	0.48

Table 2.6: Fitted values of scaling exponent α in $\chi \sim Re^\alpha$ corresponding to different ε_{L_2} thresholds in the temperature field.

Field	$\varepsilon_{L_2} = 10^{-6}$	$\varepsilon_{L_2} = 10^{-4}$	$\varepsilon_{L_2} = 10^{-2}$
θ	Runtime	Runtime	Runtime
	Memory	Memory	Memory

Figure 2.11: Evaluation of the computed exponent α : green color indicates that the MPS representation has an advantage, red color indicates that the MPS representation does not show a benefit, while yellow means that the MPS representation is comparable to FD.

In this flow regime, we expect to observe a transition from laminar to turbulent flow. Based on our scaling analysis, we note that it becomes difficult to obtain a runtime advantage for the Matrix-Vector multiplication (Matvec) operation. Furthermore, within the investigated range ($Re \leq 10^4$), we observe no saturation of the bond dimension χ ; instead, the data closely follow a power law. We have to stress that we focus our study on the instantaneous compression of DNS fields, where we limit the instantaneous error to a fixed tolerance. While this analysis showed that the required bond dimension χ continues to grow (following a power law) rather than saturating, it only represents the cost of compressing a single snapshot in time. An actual numerical simulation involves a time evolution, and it is possible that the MPS representation, even with modest χ , can still faithfully

capture the essential flow features and physics when integrated over time. Full MPS simulations are now necessary and are currently being conducted to estimate the required bond dimensions needed to accurately predict key system responses, such as the Nusselt number, which tends to be accumulated as a time average from a larger number of simulations with different initial conditions. Further investigation is required, however, to fully confirm this.

Related work by Hölscher *et al.* [8] on turbulent DJ/DT flow with periodic boundaries reports similar power-law behavior for $Re \lesssim 10^4$ demand a subsequent deviation from power law with eventual saturation of the required χ at higher Re . Accordingly, we cannot exclude that χ may ultimately saturate here as well, which could open the door to an advantage beyond the arithmetic-intensity gains discussed later for MPS.

2.1.4 Case Study III: Flow in an S-Bend



This section examines the flow in an S-bend formed by two consecutive bends, as defined in the initial research proposal's case studies. The setup is illustrated in Fig. 2.12. The comparison is between MPS and reference data generated via the finite volume method, implemented in the industry-standard HELYX toolbox (Engys) [4].

$$\mathbf{U} = \left(\frac{4u_{max}}{h^2}y(h-y), 0, 0 \right)$$




Figure 2.12: S-bend computational domain 64×16 . The fluid domain is an S-shaped duct formed by two consecutive bends. At the inlet, a laminar Poiseuille (parabolic) profile is prescribed with $u_{max} = 1$.

$N = N_x \times N_y$	32×128 (4096)	64×256 (16384)	128×512 (65536)
Re	10	100	1000

Table 2.7: Summary of grid size N with Reynolds number Re

In this study, we consider the single-bend configuration described in Deliverable 1.1 [22] and analyze the converged steady (stationary) flow over Reynolds numbers $Re \in [10, 10^3]$. The problem is discretized on three successively refined grids shown in Table 2.7—coarse (32×128), medium (64×256), and fine (128×512)—with the total number of cells $N = N_y N_x$ (where N_y and N_x denote the counts in y -direction and x -direction, respectively). The grid- Re combinations are chosen such that for every order-of-magnitude increase in Re , the grid count doubles in each spatial direction, yielding a grid-growth exponent $\beta = \log_{10}(4) \approx 0.60$. Since α and β are coupled through the grid, a different grid at each Re changes the resolved field structure and hence potentially the bond dimension, the exponents reported here are self-consistent for the present choice of grids. In future investigations, grids could be selected based on physical resolution requirements (e.g. boundary layer scaling $\delta \sim Re^{-1/2}$), and both α and β recomputed self-consistently. The velocity components u and v of the steady state are then converted into an MPS. Fig. 2.13 therefore quantifies the compressibility by the χ needed to reach a fixed fidelity on the specified grid. We expect most of the representational complexity to be localized in and around the bends (near the outer/inner walls and in the junction between bends) where curvature induces cross-stream Dean vortices, strong shear, and possibly separation; these regions are expected to drive the required MPS bond dimension χ . Upstream and far downstream the fields are smoother and should be easier to represent in the MPS format.

Scaling Analysis of S-Bend Flow

Section 2.1.1, an asymptotic runtime advantage requires $\alpha < \beta/3 \approx 0.2$, which is satisfied by all fitted values in Table 2.8 except for the u -field at $\varepsilon_F = 10^{-10}$. An asymptotic memory advantage requires $\alpha < \beta/2 \approx 0.3$, which is satisfied by all fitted values except for the u -field at $\varepsilon_F = 10^{-10}$. Based on this scaling analysis in Table 2.14, we conclude that the S-bend geometry represents a suitable candidate for future investigation.






Figure 2.13: State fidelity analysis for given fidelity threshold for different Re, on different grid size (coarse to fine grid). (a) shows u , the horizontal component of the fluid field and (b) shows v , the vertical component of the fluid field.

Field	$\varepsilon_F = 10^{-10}$	$\varepsilon_F = 10^{-8}$	$\varepsilon_F = 10^{-6}$
u	$\alpha \approx 0.315$	$\alpha \approx 0.124$	$\alpha \approx 0.122$
v	$\alpha \approx 0.165$	$\alpha \approx 0.073$	$\alpha \approx 0.081$

Table 2.8: Scaling fits $\chi \approx \text{Re}^\alpha$ at fixed ε_F .

Field	$\varepsilon_F = 10^{-10}$	$\varepsilon_F = 10^{-8}$	$\varepsilon_F = 10^{-6}$
u	Runtime	Runtime	Runtime
	Memory	Memory	Memory
v	Runtime	Runtime	Runtime
	Memory	Memory	Memory

Figure 2.14: Evaluation of the computed exponent α : green color indicates that the MPS representation has an advantage.

2.1.5 Discussion

Our analysis shows that the efficiency of the TN framework is highly problem-specific, with performance advantages varying substantially across the laminar, (transition towards quasi) turbulent, and industrial use-cases. A key finding is that the scaling of the required χ with Re depends on the chosen accuracy ε . A less strict tolerance consistently results in a more favorable power-law with a smaller slope. Notably, for all three benchmarks, the required χ follows a power-law. We have not yet observed a saturation in χ , where its growth would level off despite increasing problem complexity. However, we anticipate that in 2D simulations, such saturation is likely to occur as the flow transitions further into the (quasi) turbulent regime. Investigation of this phenomenon remains an important area of ongoing research, as the scalability of χ is the critical factor for the practical applicability of the method.

The primary obstacle to scaling modern simulations is often not raw computational power, but the prohibitive cost of data movement. Our work demonstrates that a TN approach offers a powerful solution to this memory-bandwidth bottleneck. The core of the method is to perform the entire simulation within a compressed subspace, replacing expensive global operations on data of size N with efficient local tensor contractions. This fundamentally alters the scaling of data access, making it dependent on the bond dimension χ and number of low-rank tensors $\log N$, where χ acts as a direct control for the simulation's accuracy. In practical terms, this could potentially reduce overhead between the memory and CPU (or GPU), all at a controllable approximation error set by the chosen χ . However, our main conclusion is that the theoretical advantage of this framework is ultimately governed by the physics of the flow. We have shown that for problems whose structure can be efficiently compressed into an MPS, characterized by modest growth in χ , the benefits are profound, substantially reducing memory and runtime in the scaling. Conversely, for flows nearing turbulence that demand a large χ to maintain accuracy, a memory advantage can only be achieved by relaxing the error tolerance, as demonstrated by the results in Fig. 2.11, or a saturation setting in at higher Re. However, the extent to which this relaxed tolerance impacts the physical accuracy of the flow dynamics requires further, detailed investigation. Crucially, these classical performance benchmarks provide a concrete basis for estimating the potential of quantum computation. As classical TN algorithms typically exhibit a runtime that scales cubically with χ , their quantum counterparts are expected to reduce this cost to a more favorable quadratic scaling. This allows us to establish a lower bound on the potential quantum speedup, which

grows at least linearly with the bond dimension χ .

3.1 Variational Quantum Algorithms

A promising platform to run an algorithm analogous to the TN approach for CFD is the VQA [23, 24]. In this scheme, the MPS is replaced by a parameterized quantum circuit called *ansatz* circuit, that converts a fixed, trivial input state into the desired output quantum state. Similarly, the MPO is replaced by a fixed quantum circuit. VQAs establish a hybrid quantum–classical framework ideally suited for near-term noisy quantum hardware. In a VQA the parameters of the ansatz circuits, i.e., qubit rotations, are updated iteratively by a classical optimizer that minimizes a cost function dependent on the problem, e.g., nonlinear CFD problems. There are two potential speedups that this translation offers: (i) the reduction in the scaling of the number of required variational parameters, and (ii) the algorithm realizes the equivalent to matrix-vector (MPO-MPS) multiplication as an evaluation of shallow circuit and single-qubit measurement on a quantum computer.

This section provides the theoretical foundation and initial numerical evidence supporting the potential for quantum speedups in CFD. We begin in Section 3.1.1 by introducing the fundamental concepts of VQAs, alongside a quantum version of the Nyquist-Shannon theorem [2], which provides a theoretical bound on the circuit width required to represent a function state. Moving to the core findings in Section 3.1.2, we demonstrate the potential of VQAs in representing 3D turbulent fields, showing that they require fewer parameters than the equivalent MPS representation. Furthermore, using scaling arguments specific to this case, we project that VQAs are expected to outperform MPS methods in the large system size limit, yielding significant memory and computational gains. Section 3.1.3 introduces the Tensor-Programmable VQA (TP-VQA), a method allowing the direct translation of TN algorithms into VQAs that only requires the measurement of a single ancillary qubit for cost function evaluation. We acknowledge that while fundamental scaling laws are robust, prefactors, which can be improved through better ansatz circuits are currently ignored. Therefore, the purpose of the subsequent small-scale studies (Sections 3.1.4 and 3.1.5, solving 1D PDEs) is to benchmark the required quantum resources (qubits, measurements, circuit depth) in direct analogy to the bond dimension χ . These simple case studies are designed to refine the VQA design and gather crucial engineering data before attempting large, potentially costly, simulations. Quantum circuits tend to become too deep for current noisy quantum hardware for solving complex tasks. Reducing circuit depth is therefore a key design goal. In Section 3.1.6, we discuss a strategy that trades circuit depth for circuit width, i.e., the number of circuit qubits, by using auxiliary qubits together with mid-circuit measurements and classical control to emulate the same transformation with a shallower, non-unitary circuit. Finally, we address a major challenge in Section 3.1.7, proposing an approach to mitigate the barren plateau problem, a phenomenon where the optimization landscape flattens due to the local circuit structure becoming indistinguishable from a random circuit. The section concludes with a summary of these findings and an outlook on VQA development.

3.1.1 Amplitude Encoding and Quantum Nyquist-Shannon Theorem

We first describe how a fluid field or a classical scalar field, which maps a spatial point to a relevant physical value, is mapped into a quantum state. We consider n qubits with a local Hilbert space dimension of 2 spanned by orthonormal computational basis states $|0\rangle$ and $|1\rangle$. For n qubits, the computational basis states are made up of all the possible Kronecker products of local basis states, spanning the full Hilbert space of dimension 2^n .

When one looks at individual basis state, e.g. $|0010\rangle = |0\rangle \otimes |0\rangle \otimes |1\rangle \otimes |0\rangle$, one notices that it forms a bit string that is unique for each computational basis state, and it represents a number, in a binary form, between 0 and $2^n - 1$. The amplitude encoding gives a one-to-one relation between the basis state and each point in the discretized grid. For a one-dimensional field defined between an interval a and b , the natural approach is to have an equal interval grid on points $x_i = a + \frac{b-a}{2^n}i$ ($i \in \{0, 1, \dots, 2^n - 1\}$), such that field values $f(x_i)$ are encoded as an amplitude of the basis state $|i_2\rangle$, where i_2 denotes the bit string corresponding to the binary representation of i . The resulting state is $|f\rangle \propto \sum_{i=0}^{2^n-1} f(x_i) |i\rangle$ with the constant of proportionality that ensures the inner product to be 1, $\langle f|f\rangle = 1$. Naturally, the individual qubit controls the features at different length scales, for this deliverable, we set a convention that the first qubit controls the largest digit and the n -th qubit the smallest. Finally, the encoding naturally extends to encoding fields on higher dimensions, by discretizing each dimension independently as described. For a square lattice, the total number extends to dn , where d is the physical dimension of which the field resides.

We now state the quantum Nyquist-Shannon theorem and show the sketch of the proof. The theorem states that in order to amplitude-encode a field that oscillates with wave number $\frac{2\pi}{\lambda_c}$, one needs the n -th qubit to represent






Figure 3.15: Linear quantum circuit and its parameterization. a. Circuit diagram of an initial product state $|0\rangle^{\otimes n}$ evolving via local interactions, depicted by the red vertical lines that connect the world lines of the qubits, towards the output state $|\psi\rangle_{\text{out}}$. b. For performing VQA, we define a block of local interaction is parameterized by local rotation gates ($U(\theta_i)$ for i -th parameter) and fixed two-qubit gate (the blue vertical line).

the length scale of λ_c/π . If we set the computational box size to unity ($a = 0$ and $b = 1$), one needs at least $q = \log_2 \pi/\lambda$ qubits in order to fully resolve the function.

This statement can be proven by considering the thermodynamic limit ($\lim_{n \rightarrow \infty}$) of the explicit tensor network representation of the function $f(x) = \sin(2\pi/\lambda)$. Inductively, the density matrix ρ_i of qubit i can be shown to have the following form

$$\rho_i = \begin{pmatrix} \frac{2k - \sec(2^{-i-1}k) \sin((2+2^{-i-1})k) + \tan 2^{-i-1}k}{4k - 2 \sin 2k} & \frac{2k \cos 2^{-i-1}k - \frac{\sin 2k}{\cos 2^{-i-1}k}}{(4k - 2 \sin 2k)} \\ \frac{2k \cos 2^{-i-1}k - \frac{\sin 2k}{\cos 2^{-i-1}k}}{(4k - 2 \sin 2k)} & \frac{2k - \sec(2^{-i-1}k) \sin((2-2^{-i-1})k) - \tan 2^{-i-1}k}{4k - 2 \sin 2k} \end{pmatrix}, \quad (3.6)$$

where $k = 2\pi/\lambda$. From the expression, for $2^{-i-1}k < 1$, every element of the reduced density operator converges exponentially towards $1/2$ with i , which is equivalent to a convergence to $|+\rangle \langle +| \propto |0\rangle \langle 0| + |0\rangle \langle 1| + |1\rangle \langle 0| + |1\rangle \langle 1|$ state. By rearranging the inequality we obtain the threshold number of qubits for resolution, $q_c(\lambda) = \log_2 \pi/\lambda$, given in the main text. Thus, beyond $i = \lceil q_c \rceil$, the qubit state becomes independent of the function's oscillations, meaning qubit q_c is the last qubit that contributes towards the encoding. For more detailed discussion and proof, cf. [2].

The quantum Nyquist-Shannon theorem immediately fixes the width of the circuit. For the fluid dynamical simulations, the smallest length scale of resolution is given by the Kolmogorov length scale $\xi_K < 1$, which scales like $O(\nu^{3/4})$ with the viscosity ν , indicating that the circuit width w grows only like $O(\log_2 \frac{b-a}{\xi_K})$ in the turbulent regime where ξ_K is small in comparison to the box size. Furthermore, we estimate the required depth of the circuit for capturing the key properties of the fields. Capturing κ -point correlators, represented by the structure functions, requires, at worst $O(\kappa^{10}w^2)$ variational parameters [25] for nearest-neighbor variational ansatz circuit drawn in Fig. 3.15, for which the depth is given by w . In the next section, we provide numerical evidence that in QCFD applications, this bound is hardly ever reached, but a more practical bound is derived by analyzing the quantum mechanical properties of the Navier-Stokes equations.

The aforementioned $O(\kappa^{10}w^2)$ scaling, however, is not practical as the coefficient of the scaling tends to be very large and we rarely encounter structureless, pure noise-like fluid field features as a result of a PDE-based time evolution. Therefore, in order to find a practical bound, it is important to look at structure of a PDE in the context of quantum mechanics. We find that time evolution of a nonlinear PDE that is analogous to the Navier-Stokes equation, in Fourier space, has a structure such that correlations spread only linearly throughout the qubits. Therefore, the $O(w^2)$ (consequently $O((\log_2 \frac{b-a}{\xi_K})^2)$ in the turbulent regime) bound holds with a much smaller coefficient in the CFD context, where a detailed discussion is presented in our future publication [2].

In what follows, the optimization and VQA are performed in the real space instead of in Fourier space. This is because complications introduced by boundary conditions and initial state preparation in Fourier space are avoided in the real space representation, thus making the practical implementation of a quantum variational PDE solver possible. Nevertheless, the above scaling arguments still hold in real space, as the quantum Fourier transform operation does not generate much entanglement as discussed in [26], so the effect will be at worst adding $O(1)$ depth to the ansatz circuit for large enough n .

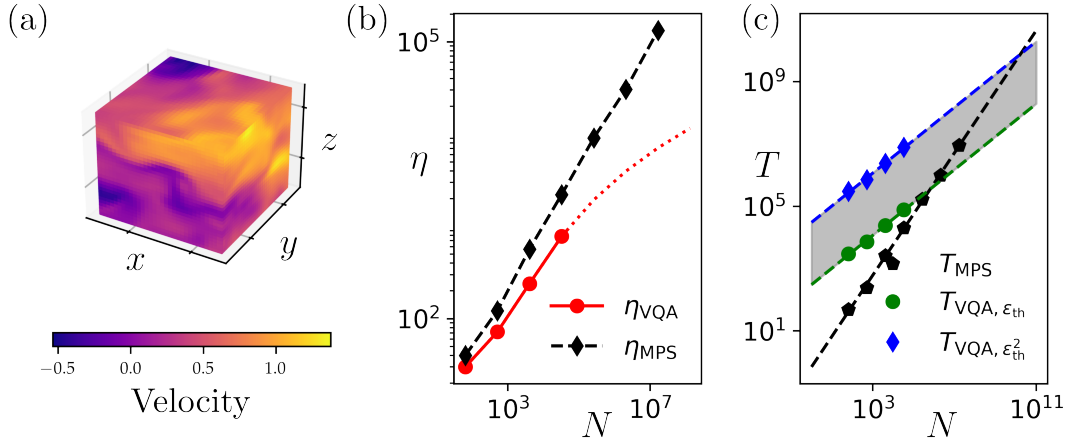


Figure 3.16: (a) Snapshot of y -velocity component of the Turbulent flow data [27, 28, 29] used for the scaling analysis ($N = 32 \times 32 \times 32$). (b) Scaling of number of parameters required to describe the field on N grid points for obtaining a faithful representation of a state below a fixed error threshold $\epsilon_{th} = 0.01$. The red dotted line is an order 3 polynomial fitted to the η_{VQA} data points, and the black dashed line is a linear fit in log-log scale for the η_{MPS} data points. The scaling we observed for VQA is preliminary and further research is required. (c) The scaling of the bounds on required quantum resources. The green line is for the lower bound ($T_{QC, \epsilon_{th}}$) and the blue line is for the upper bound (T_{QC, ϵ_{th}^2}) with $C_{\epsilon_{th}} = C_{\epsilon_{th}^2} = 1$. These are plotted against the computational cost for MPS. They are derived from (b) by fixing the error tolerance to $\epsilon_{th} = 0.01$. The dotted lines are linear fits in log-log scale for each cost. For this set of constants of proportionality ($C_{\epsilon_{th}}$ and $C_{\epsilon_{th}^2}$), the true computational cost lies within the gray area. Data taken from [30]

3.1.2 VQA vs TN: Scaling Analysis

We confirm the theoretical scaling derived in Section 3.1.1 by comparing the number of parameters η required for faithfully expressing the turbulent flow. The turbulent flow is chosen as it is the worst possible structure that may emerge as a result of CFD simulations in terms of complexity (Fig. 3.16 (a)). The infidelity, ϵ_F , is used to quantify the encoding error between the true state ($|\psi_{true}\rangle$) and the approximate state ($|\psi_{approx}\rangle$) obtained variationally with an MPS or quantum circuit. The definition of the infidelity metric is introduced in Section 1.3. Furthermore, for the quantum representation we use the same ansatz circuit as in the previous subsection (see Fig. 3.15), and use an ADAM optimizer with finely tuned parameters until the error becomes smaller than a threshold $\epsilon < \epsilon_{th}$.

In Fig. 3.16 (b), we compare the number of parameters, η , required for MPS and VQA. For this example, we find that VQA requires a smaller number of parameters for all grid sizes, i.e. the number of grid points, N investigated. Furthermore, the difference $\eta_{MPS} - \eta_{VQA}$ monotonically increases with N , suggesting that VQA may have a scaling advantage in the number of parameters over its classical counterpart, namely, the MPS method. Here we also find the scaling $\eta_{MPS} \propto N^{0.65}$ and $\eta_{VQA} \propto N^{0.52}$, where VQA has a better scaling than MPS. We have also found that a polynomial of order 3 in $\log_2 N$ fits well (red dotted), however, due to the limitations in the achievable number of qubits, the true scaling is inconclusive and requires further research.

Next, we compare the computational cost, T , of the MPS method to VQA. As explained in the previous part, the computational cost of MPS scales as $T_{MPS} = \chi^3 \log_2 N$, where χ is the bond dimension. For a fixed threshold ϵ_{th} , however, χ may depend on N , since increasing the system size can introduce non-trivial structures that must be resolved. Therefore, in Fig. 3.16 (c), we take the values of $\chi_{<0.01}$, the minimum chi required for representing the field with error below $\epsilon_{th} = 0.01$ obtained in (b), and plot $T_{MPS} = \chi_{<0.01}^3 \log_2 N$. Here we find a scaling $T_{MPS} \propto N$ (c, black dashed).

For VQA, we define the cost T_{VQA} as the number of measurements required to compute the gradients of the cost function for all parameters down to the precision $\epsilon_{th} = 0.01$. Using the Hadamard test [23], the best-case cost is $T_{VQA, \epsilon} = C_{\epsilon_{th}} \eta^{VQA} / \epsilon_{th}$ for evaluating the η^{VQA} gradients. This occurs when the ancillary qubit's Bloch vector aligns completely with the measurement basis. Conversely, the worst-case cost scales as $T_{VQA, \epsilon_{th}^2} = C_{\epsilon_{th}^2} \eta_{VQA} / \epsilon_{th}^2$, where the $1/\epsilon_{th}^2$ factor arises from the Monte Carlo sampling limit. In our example, we find that both $T_{VQA, \epsilon_{th}}$ and T_{VQA, ϵ_{th}^2} scale as $\propto \sqrt{N}$ (or potentially $O((\log N)^3)$), which is more favorable than the $T_{MPS} \propto N$ observed for MPS. Note that sampling can be parallelized trivially. In principle, this approach would impact the constant prefactors of the scaling, rather than altering the fundamental asymptotic scaling itself.

In conclusion, our scaling analysis suggests a potential for quantum advantage in CFD. The VQA method



demonstrates more efficient scaling, with a computational cost of $O(\sqrt{N})$ (and potentially $O((\log N)^3)$) compared to the $O(N)$ scaling of the classical MPS method. However, the observed scaling for the number of parameters, and hence the computational costs disagrees with the theoretical scaling provided in Section 3.1.1. This may be blamed on the barren-plateau phenomenon [31] that prohibits the trainability of the quantum circuit, and a suboptimal choice of the ansatz circuit structure. These problems are tackled in Section 3.1.7.

Nevertheless, this suggests that as problem sizes increase, the performance gap between quantum and classical approaches could widen. However, the VQA analysis is based on system sizes that are reachable with currently available classical computational resources. Focusing on practical implementations on quantum hardware to account for noise and other architectural limitations is also a potential future direction. This observed initial theoretical advantage provides strong motivation for continued research.

3.1.3 Tensor-Programmable VQA

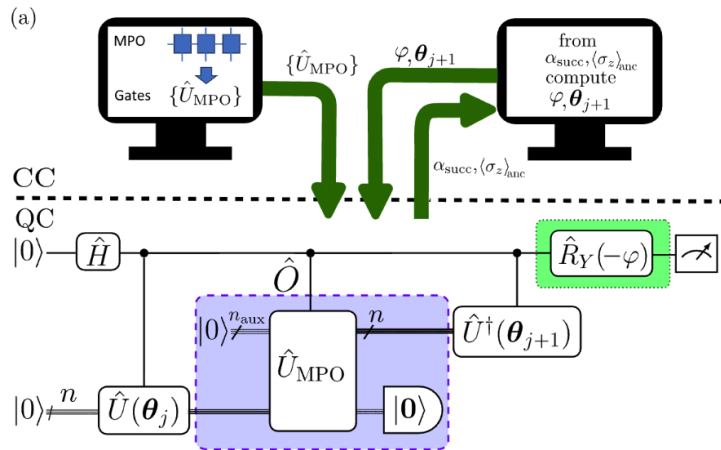


Figure 3.17: Sketch of the hybrid classical-quantum strategy using MPO-programmed unitaries and postselection in a VQA setting. The figure is taken from [30].

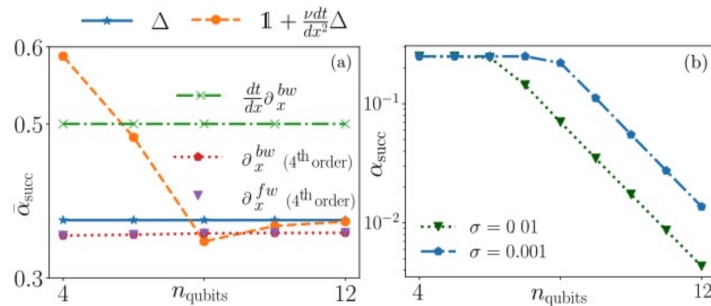


Figure 3.18: Success probabilities vs. system size. (a) Empirical mean success probability $\bar{\alpha}_{\text{succ}}$ for linear stencil operators: the Laplacian Δ (blue) and first-derivative schemes (green, red, purple) remain essentially size-independent with $\bar{\alpha}_{\text{succ}} \approx 0.35$ – 0.50 , while the implicit diffusion step $1 + \nu \frac{\Delta t}{\Delta x^2} \Delta$ (orange) decreases with n_{qubits} before flattening. (b) Nonlinear pointwise multiplication: the success amplitude α_{succ} stays near-constant until the feature (Gaussian width σ) is resolved, after which it drops with n_{qubits} ; narrower features ($\sigma = 0.01$) decay earlier/faster than wider ones ($\sigma = 0.001$). The figure is taken from [30].

In the previous sections, we established the correspondence between a parametrized quantum circuits and an MPS: the entries of MPS tensors play the role of variational parameters, while in a VQA the analogous degrees of freedom are the gate parameters (e.g., rotation angles/phases) of the constituent unitaries. An analogous link can be drawn for operators: matrices (e.g., differential or linearized CFD operators) acting on a state vector can be represented by quantum circuits. The main complication is that some CFD operators are non-unitary, whereas quantum circuits natively require unitary transformations. Two established solutions are (i) block-encoding, which embeds a non-unitary operator into a larger unitary acting on system plus ancilla [32], and (ii) variational unitary

MPOs [33], in which we optimize a (variational) unitary MPO that approximates the target MPO. Both approaches require ancillary qubits and typically rely on measurements of them to identify the successful applications, yielding a probabilistic realization of the desired operator. Thus, there is a natural analogy: MPS \leftrightarrow variational state ansatz, and MPO \leftrightarrow probabilistic unitary. This strategy—using tensor networks as a programming paradigm for quantum algorithms—was already motivated by Lubasch et al. [23].

The required circuit resources follow directly from the MPO structure. In particular, the gate width in a unitary-MPO realization is set by the MPO bond dimension χ_o : each entangling block in the unitary-MPO circuit typically acts on $m \gtrsim \lceil \log_2 \chi \rceil$ qubits. Thus, larger χ_o implies larger multi-qubit gates or, at fixed width, deeper decompositions (more layers with smaller multi-qubit gates). As the operator implementation becomes probabilistic, with success probability p potentially below unity, this probability directly determines the number of measurement shots required to reach a desired statistical accuracy, since only successful applications contribute to the cost functional estimate.

We assessed the impact of these factors within a VQA setting in our recent preprint [30]. For the linear operators studied, the average success probability remained constant or even converged with system size (see Fig. 3.18a). For nonlinear pointwise multiplication, the success rate stayed roughly constant until the relevant features were resolved, after which it exhibited a decrease (see Fig. 3.18b). Crucially, in neither case did the overall VQA algorithm exhibit an exponential decay of success probability with the iteration count.

3.1.4 Case Study IV: Nonlinear Schrödinger equation



In this section, we discuss the scaling of the number of qubits, the number of CNOT gates, and the circuit depth required to variationally solve the nonlinear Schrödinger equation (NLSE). In our case, we focus on the quadratic nonlinearity, as it represents one of the most fundamental and physically relevant forms of nonlinearity encountered in diverse quantum systems. The insights gained from this analysis will also guide the design and resource estimation of variational circuits for other models exhibiting similar quadratic nonlinearities. This focus allows us to establish a concrete benchmark for scaling behavior and algorithmic performance in representative nonlinear quantum problems.

The nonlinear Schrödinger equation (NLSE) is a one-dimensional time-independent equation that helps model various phenomena such as the dynamics of light in nonlinear optics [34, 35, 36], envelope solitons and modulation instabilities in plasma physics and surface gravity waves [37], and characteristics such as superfluidity and vortex formation in Bose-Einstein condensates (BEC) [38, 39, 40, 41, 42], to name a few. In dimensionless form, the NLSE is given as

$$\left[-\frac{1}{2} \frac{d^2}{dx^2} + V_0(x - x_0)^2 + g|f(x)|^2 \right] f(x) = Ef(x). \quad (3.7)$$

Here, $f(x)$, with x being spatial coordinates, represents a normalized single real-valued function defined over the interval $[a, b]$. g denotes the strength of the nonlinearity, and V_0 is the depth of the quadratic potential $V(x)$ centered around $x_0 = \frac{b-a}{2}$. Furthermore, we consider periodic boundary conditions such that $f(b) = f(a)$ and $V(b) = V(a)$. It is worth highlighting that we employ the classical method of the imaginary-time evolution [43, 40, 44] to benchmark the variational results.

For the variational ground state problem of the NLSE, we discretize the spatial domain and consider finite-difference method. The cost function is defined as the sum of the kinetic, potential, and interaction energy, $\langle\langle E \rangle\rangle = \langle\langle E_P \rangle\rangle + \langle\langle E_I \rangle\rangle + \langle\langle E_K \rangle\rangle$ [23, 45], where

$$\begin{aligned} \langle\langle E_P \rangle\rangle &= \sum_{k=0}^{N-1} |\psi_k|^2 V_k, & \langle\langle E_I \rangle\rangle &= \sum_{k=0}^{N-1} \frac{g}{\delta} |\psi_k|^4, \\ \langle\langle E_K \rangle\rangle &= -\frac{1}{2\delta^2} \sum_{k=0}^{N-1} (\psi_k^* \psi_{k+1} - 2\psi_k^* \psi_k + \psi_k^* \psi_{k-1}), \end{aligned} \quad (3.8)$$

and the minimum value of the cost function represents the ground state solution. Here, $\psi_k = \sqrt{\delta} f(x_k)$ such that



the normalization condition $1 = \delta \sum_{k=0}^{N-1} |f(x_k)|^2 = \sum_{k=0}^{N-1} |\psi_k|^2$ is satisfied. Each term in the cost function requires a separate quantum circuit (refer to [23, 45] and [46] for two different circuit constructions). It turns out that the maximum number of qubits and gate resources required to evaluate the cost function is determined by the nonlinear interaction energy term $\langle\langle E_I \rangle\rangle$. While the mapping of the trial solution requires a register of $\log_2(N)$ qubits, the encoding of the nonlinearity term demands $3 \log_2(N)$. The need for three quantum registers to encode the nonlinearity reflects the implications of the no-cloning theorem in quantum mechanics, which necessitates the preparation of multiple copies of the trial solution.

Qubit scaling vs problem size

The number of qubits n scales as $3 \log_2(N) + 1$, where N is the number of grid points of the problem. Since the number of grid points ($N = 2^{\frac{n-1}{3}}$) grows faster than the qubit requirements, the algorithm does not face inherent scalability limitations. However, ordinary classical computers can typically handle simulations of up to about 16–20 qubits ($N = 2^6 - 2^7$ grid points), while modern supercomputers, depending on architecture and memory, can extend this limit to approximately 40–45 qubits ($N = 2^{14} - 2^{15}$ grid points). On the contrary, current NISQ devices already offer more than 1000 qubits, which in principle would be sufficient to represent problems defined on grids as large as 2^{300} points. However, the practical solution of such problems remains currently out of reach due to noise, limited circuit depth, and fidelity constraints.

Circuit depth

Within gate-level simulations, we analyze the circuit depth requirements to capture the variational ground state of the NLSE with up to 99% fidelity. We systematically analyze the number of two-qubit gates and circuit depth required to achieve $> 99\%$ fidelity for $n = 7$ ($N = 4$), $n = 10$ ($N = 8$) and $n = 13$ ($N = 16$) [see [46] for more details] and extrapolated it to relatively bigger system sizes in the Fig. 3.19.




Figure 3.19: Scaling of number of (a) CNOT gates and (b) circuit depth with respect to problem size in terms of number of qubits n . Here, black, blue, and purple colors indicate the weak, moderate, and strong regimes of nonlinearity.

3.1.5 Case Study V: Burgers' equation



In this section, we highlight the scaling of the number of qubits and the number of CNOT gates required to efficiently solve the nonlinear dynamics of the one-dimensional viscous Burgers' equation. Our analysis shows a three-fold reduction in the number of two-qubit gates compared to the conventional implementation of quantum circuits, indicating a remarkable improvement in circuit efficiency while maintaining high-fidelity simulation of the underlying dynamics. Studying this model is particularly important as the Burgers' equation serves as a prototypical nonlinear partial differential equation, capturing essential features of turbulence and shock formation.



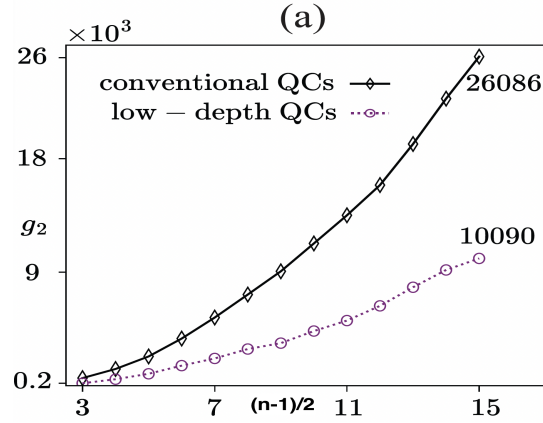


Figure 3.20: Scaling of two-qubit gate count (g_2) with respect to problem size n in terms of number of qubits. Here, black, and purple colors indicate the conventional and low-depth implementations of quantum circuits. This plot was taken from [48].

The insights gained from this analysis provide a valuable foundation for extending VQAs to more complex fluid dynamical systems governed by nonlinearities of similar nature.

Probing the nonlinear dynamics of fluid configurations is essential for advancing fluid dynamics simulations, as it enables deeper insight into turbulence, shock formation, and other complex phenomena that govern real-world transport and flow processes. Here, we consider the time-dependent one-dimensional viscous Burgers' equation [47, 23, 24]

$$\partial_t u(x, t) = [\nu \partial_x^2 - u(x, t) \partial_x] u(x, t) . \quad (3.9)$$

Here, the first term on the right side of the equation represents diffusion, where ν denotes the coefficient of kinematic viscosity. The second term corresponds to convection, which is inherently nonlinear in nature. ∂_x and ∂_t denote the spatial and temporal derivatives, respectively, while $u(x, t)$ is the real-valued fluid field, and $x(t)$ represents space (time). Eq. (3.9) models various flow regimes, including laminar ($\nu \geq 1$) and turbulent ($\nu < 1$) behavior, and small instances of the problem can be solved on a classical computer using a range of straightforward techniques, such as DNS, FD methods, finite volume methods, spectral methods, and other approaches. Quantum algorithms for simulating nonlinear Burgers dynamics have been proposed in [23, 24] and subsequently subjected to further theoretical analysis and refinement in [46, 48].

The cost function for the Burgers dynamics is defined as the squared residual of the variational state and the time-evolved state and is given as,

$$\mathcal{C} = ||u_{t+\tau} - [1 + \nu \partial_x^2 - |u_t\rangle \partial_x] |u_t\rangle|^2 . \quad (3.10)$$

For each time step t , the cost function defined in Eq. (3.10) is minimized, and long-time dynamics can be investigated by iteratively repeating this procedure. In [46, 48], we analyzed the expressibility of various parameterized quantum ansätze, the convergence properties of different optimization strategies, and the reduction of gate counts in quantum circuits. These studies provide quantitative benchmarks for assessing the efficiency and scalability of variational approaches to nonlinear dynamical problems. Moreover, they highlight the trade-offs between circuit depth, optimization accuracy, and hardware feasibility, thereby offering valuable guidance for the design of resource-efficient quantum algorithms.

Qubit scaling vs problem size

For the Burgers equation, nonlinearity is a convection-like term, rather than a quadratic term. Therefore, the number of qubits n scales as $2\log_2(N) + 1$, where N is the number of grid points of the problem. Gate-level simulations of Burgers dynamics can be carried out on classical computers for system sizes up to $N = 2^9$ grid points, while state-of-the-art supercomputers can extend this limit to approximately $N = 2^{22}$. However, since dynamical simulations require the solution to be computed at every time step, the computational cost escalates rapidly and becomes infeasible for large systems. Consequently, only small-scale gate-level VQA simulations with $N \approx 2^{10} - 2^{12}$ grid points can be simulated on supercomputers within a reasonable timeframe.



Two-qubit gate count

Focusing on the turbulent regime of fluid flow, we analyze the two-qubit gate requirements necessary to capture shockwave behavior with fidelities exceeding 99%. To this end, we proposed and investigated low-depth implementations of quantum circuits for simulating Burgers' dynamics, demonstrating a three-fold reduction in the number of two-qubit gates compared to conventional approaches. This reduction not only improves scalability on near-term quantum hardware but also mitigates the impact of noise, as two-qubit gates are typically the dominant source of error. Furthermore, our results establish that variational representations of nonlinear fluid phenomena can be achieved with significantly shallower circuits than previously anticipated, thereby advancing the feasibility of practical quantum simulations of turbulence, in current NISQ devices. In the Fig. 3.20, we illustrate how the two-qubit gate count scales with increasing system size. In the context of fault-tolerant quantum computation, performance assessments are typically characterized in terms of T-gate counts, which serve as a primary metric of algorithmic cost. Consequently, an extensive resource analysis should be undertaken to estimate and predict the requirements for implementing QCFD within error-corrected quantum architectures, enabling realistic projections for large-scale, fault-tolerant deployments.

3.1.6 Case Study VI: Circuit Optimization



Quantum algorithms often yield circuits with considerable depth, especially as the complexity of the computational task increases. These deep circuits typically involve highly entangled states that must retain coherence for extended durations, which presents a major challenge for current quantum hardware. Consequently, reducing circuit depth has become a central goal in quantum algorithm design, with a substantial body of research dedicated to this problem. One way a shallower depth can be achieved is by increasing the number of qubits, which in turn often requires the use of mid-circuit measurements and classically controlled operations. In practice, a unitary circuit is replaced by a non-unitary circuit of a lower depth that induces effectively the same transformation on the relevant (register) qubits.

In the case of VQAs, this substitution is relevant in the employed ansatz circuits. These circuits typically consist of alternating layers of single qubit rotations and entangling two-qubit gates, with the parameters updated iteratively by a classical optimizer to minimize a cost function similar to the one discussed above for Burger's dynamics. A general ansatz circuit is presented in Fig. 3.21(a). In Fig. 3.21(b) we present a typical configuration that is used in ansatz circuits, that we refer to as *core* circuit, and in Fig. 3.21(c) its corresponding non-unitary version with a constant circuit depth. Five qubits are used for illustration purposes, however the construction is modular and can be straightforwardly extended to any number of qubits. In Fig. 3.22 we present the comparative analysis between the unitary and non-unitary ansatz circuits. Based on these results, the non-unitary construction appears to be preferable for quantum platforms that present relatively low two-qubit error rates and relatively high idling error rates. This non-unitary construction is not limited to the core circuit shown in this report but can be applied to a larger family of core circuits as long as they are constructed by a series of two-qubit gates where one starts at the point the previous ends forming a structure.

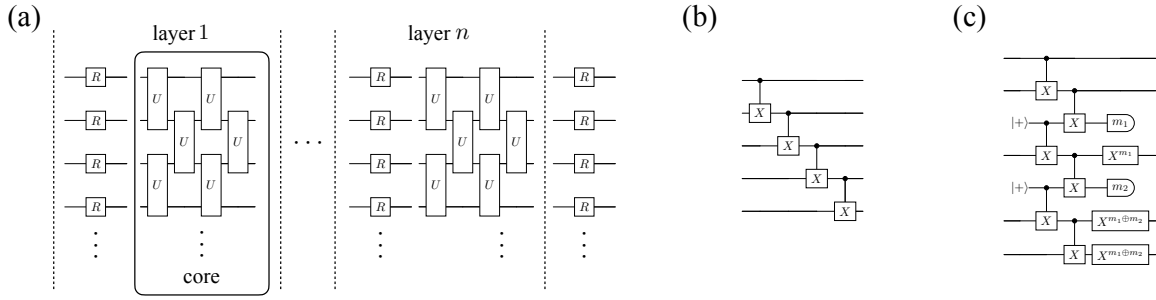


Figure 3.21: In panel (a) a generic ansatz circuit is presented, where the configuration of two-qubit gates for each layer is referred to as core. In panel (b) a typical core circuit is shown, which is a unitary process. In panel (c) a non-unitary process is presented that induces the same transformation as the circuit in panel (b).

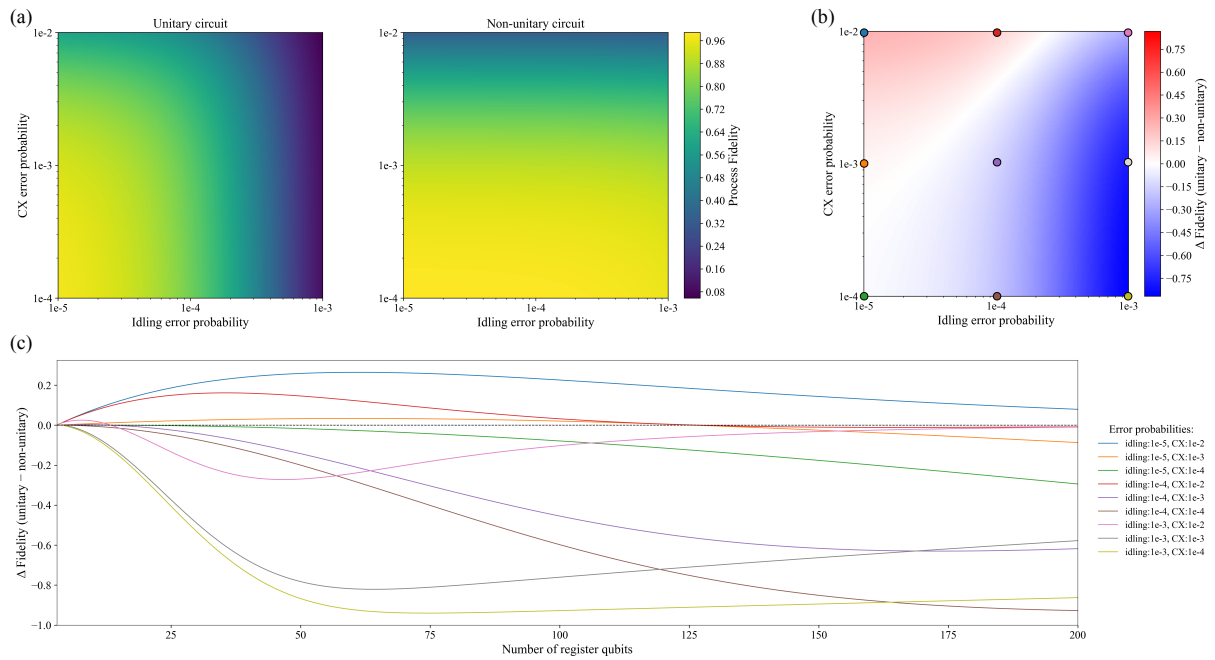


Figure 3.22: In panel (a) two plots are presented for the unitary and the non-unitary version of 50-qubit core circuit shown in Fig. 3.21(b). Both have the same range in terms of idling error probability and CX error probability: $1e-5$ to $1e-3$ and $1e-4$ to $1e-2$, respectively. The colorbar corresponds to both plots and indicates the lower bound of process fidelity of the entire circuit, which appears to be at its highest for the unitary circuit when idling error probability is low, and for the non-unitary circuit when CX error probability is low. In panel (b) the difference of the process fidelities is shown, Δ Fidelity, where positive (red) indicates an advantage for the unitary circuit and negative (blue) for the non-unitary one. Finally, in panel (c) we present the Δ Fidelity against the number of register qubits for nine different error probabilities, which are also indicated as colored dots on the plot of panel (b).

3.1.7 Quantum Sparse Optimization Near the Edge of Chaos

In performing VQA, for the variational circuit to contain the solution, it must be sufficiently flexible to explore a wide range of possible solutions (expressibility), while also ensuring the parameters are not redundant, which would lead to overparameterization. In VQAs, therefore, one must consider the trade-off between the expressibility of an ansatz and the trainability of its parameters. A parameterized quantum circuit that is too shallow may not have the capacity to represent the solution within its accessible manifold. Conversely, a deep circuit can lead to a locally random structure that causes a barren plateau, a parameter regime where the cost function gradient vanishes exponentially with the number of qubits, making the optimization untrainable (see Fig. 3.23a for schematics on circuit depth, correlation spreading, overparameterization, and emergence of a barren plateau). Ideally, then, one wants an ansatz that is deep enough to represent the solution yet structured enough to avoid the random behavior




Figure 3.23: (a) Schematic of the Training Regimes. This diagram illustrates the relationship between circuit depth (D), the causal lightcone of parameters, and the onset of the barren plateau. The ansatz is a brick-layer circuit of single-qubit rotations and two-qubit entangling gates (red blocks) as presented in Fig. 3.15. A change to a parameter early in the circuit creates an expanding lightcone of influence (yellow). As the circuit deepens, these lightcones overlap and information scrambles across an exponentially large Hilbert space (dark red region), causing parameter gradients to vanish. Our TEE-based regularizer effectively halts this spread, keeping the circuit’s dynamics at the trainable “edge of chaos” (b) Energy error ΔE relative to the exact ground state energy versus optimization steps. The standard VQE optimization (black solid line) stagnates in a barren plateau, while the TEE-regularized approach (red dashed line) successfully converges, achieving a final error five orders of magnitude smaller. (c) The value of the regularization term C_{I_3} during the regularized training. The regularizer is active for the first 500 steps, confining the state to a trainable, “non-chaotic” regime. After it is turned off, C_{I_3} increases as the state develops the necessary entanglement, showing the optimization successfully approaches the solution from the “edge” of the chaotic region. The panels are recreated based on [2] with additional data obtained for panels (b) and (c).

that impedes successful training on a large number of qubits.

We tackle this trade-off between expressibility and trainability by adapting regularization techniques from classical optimization and concepts from the theory of information propagation in lattice systems, a topic well-studied in quantum chaos and condensed matter physics. We propose a quantum regularization strategy that introduces a notion of quantum sparseness. This is achieved by penalizing parameter regimes that drive the circuit toward chaotic dynamics while also avoiding the trivial, shallow-circuit regime where the output state is classically simulable with tensor networks. The regularizer, therefore, guides the optimization to operate at the critical point—the *edge of chaos*—between the classically tractable and the chaotically intractable regimes.

This regularizer is based on a quantity called Topological Entanglement Entropy (TEE). TEE quantifies a form of many-body entanglement that is stored non-locally across the system, rather than between specific pairs of qubits. In a state with high TEE, information about the global properties of the system cannot be recovered through purely local measurements. For a system partitioned into three contiguous regions A , B , and C , it is defined as the tripartite mutual information:

$$I_3(A : B : C) = S_A + S_B + S_C - S_{AB} - S_{AC} - S_{BC} + S_{ABC}$$

where S_X is an entropy measure. For practical and experimental relevance, we chose the Rényi-2 entropy $S_X = -\log \text{Tr}(\rho_X^2)$ for the rest of this section, where ρ_X is the reduced density matrix for region X . This specific formulation cancels out local, area-law contributions to the entanglement entropy that scale with the boundary length of the regions, thereby isolating the entanglement across multiple (more than 2) qubits. Therefore, minimizing TEE acts to restrict the causal lightcone of the variational parameters. This localization of each parameter’s influence prevents the global information scrambling responsible for the barren plateau problem, thereby improving the circuit’s trainability.

To demonstrate the effectiveness of the TEE-based regularizer, here we use a quantum mechanical model, namely the nearest-neighbor Heisenberg model, and use VQA to find its ground state. The model is an ideal test case for validating the TEE-based regularization as it is a foundational model in quantum magnetism and its ground state at different parameters exhibits a rich spectrum of entanglement structures that emerge from a locally interacting quantum system. This non-trivial entanglement is not only relevant across various fields in quantum physics but also models the emergent entanglement structure of amplitude-encoded fluid fields at intermediate times resulting from short-ranged-Hamiltonian-like time evolution, as discussed in Section 3.1.1.

We determine the ground state of the one-dimensional quantum Heisenberg model on a chain of $n = 9$ qubits with periodic boundary conditions. The system’s Hamiltonian is given by:

$$H = J \sum_{\langle i,j \rangle} (\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y + \sigma_i^z \sigma_j^z) + h \sum_{i=0}^{n-1} \sigma_i^z,$$

where σ^z is a Pauli-z operator with eigenstates $|0\rangle$ and $|1\rangle$ with corresponding eigenvalues $+1$ and -1 , and σ^x is a Pauli-X operator $\sigma^x = |0\rangle\langle 1| + |1\rangle\langle 0|$. We employ a linear ansatz, consisting of local parameterized rotations and nearest-neighbor entangling gates (Fig. 3.15), to prepare the variational state $|\psi(\boldsymbol{\theta})\rangle$. The parameters $\boldsymbol{\theta}$ are then optimized to minimize a composite cost function:

$$\mathcal{C} = \langle \psi(\boldsymbol{\theta}) | H | \psi(\boldsymbol{\theta}) \rangle + \gamma \mathcal{C}_{I_3}$$

The regularization term, \mathcal{C}_{I_3} , is defined as the sum of local TEE values across the chain

$$\mathcal{C}_{I_3} = \sum_{i=0}^{n-1} |I_3(i, 2)|$$

Here, $I_3(i, 2)$ is a shorthand for the tripartite mutual information $I_3(A : B : C)$ where the regions are three contiguous single-qubit sites: $A = \{i, i + 1\}$, $B = \{i + 2, i + 3\}$, and $C = \{i + 3, i + 4\}$, with indices taken modulo n and γ is the Lagrange multiplier that penalizes the spread of the lightcone.

Fig. 3.23b. shows the training results for the model with parameters $J = 1$ and $h/J = 1$. We compare the standard optimization (black solid line) with our TEE-regularized approach (red dashed line). For the regularized training, we introduce the cost function term $\gamma \mathcal{C}_{I_3}$ and set the regularization strength to $\gamma = 0.01$ for the first 500 optimization steps. Afterward, the regularizer is turned off ($\gamma = 0$) to allow the variational state to develop the modest entanglement required by the true ground state. As shown, the regularized approach yields a five-orders-of-magnitude improvement in the final energy error, ΔE , which is the difference between the final variational energy and the exact ground state energy.

Furthermore, Fig. 3.23c. illustrates the behavior of the regularization term \mathcal{C}_{I_3} during the optimization. For the regularized run, \mathcal{C}_{I_3} is initially suppressed and then increases sharply after the regularizer is turned off at step 500. This behavior indicates that the optimization path approaches the optimum from the “non-chaotic” side of the parameter space. We conclude that the regularizer successfully confines the optimization to a region of the manifold with non-vanishing gradients. This enables convergence to a high-quality minimum that is otherwise inaccessible from the chaotic regime where barren plateaus is present.

These results suggest that directly controlling the non-local entanglement structure of the variational ansatz is a powerful strategy for improving trainability. Such physics-informed regularization strategies offer a promising path toward making VQAs practical and effective on near-term quantum hardware. Furthermore, similarly promising results have been obtained for the variational amplitude encoding of turbulent data, which are detailed in a preprint [2]. Future research should therefore analyze the resource requirements and error propagation of this method in more realistic, noisy settings.

3.1.8 Discussion

In this section, we established a quantum framework for simulating canonical CFD problems and nonlinear partial differential equations by translating classical TN methods into VQAs. By mapping the MPS ansatz to a VQA, the solution can be encoded more efficiently, requiring substantially fewer parameters than the classical method. The quantum Nyquist-Shannon theorem provides a clear guideline for qubit requirements, and the TP-VQA approach allows direct implementation of tensor network algorithms using ancillary qubit measurements. Scaling analysis demonstrate that VQAs offer favorable computational scaling of $O(\sqrt{N})$, compared to the $O(N)$ scaling of classical MPS algorithms. We also motivate the necessity of benchmarking VQAs on simple 1D problems to quantify the required quantum resources (such as qubits and measurements) on actual quantum hardware and to minimize the scaling prefactors through circuit refinement. Furthermore, several strategies for practical quantum implementation are discussed. Shallow, optimized non-unitary circuits reduce circuit depth and mitigate errors, while case studies, including the nonlinear Schrödinger equation and Burgers’ equation, confirm that VQAs can efficiently handle nonlinearities and achieve high-fidelity solutions with reasonable qubit and gate scaling. Finally, we introduce and apply a novel method designed to mitigate the challenge posed by barren plateaus. We provide an illustrative example where this method is used in a VQA to find the ground state of the one-dimensional quantum Heisenberg model. These results provide a roadmap for leveraging near-term devices to solve nonlinear problems.

4.1 Conclusion

This report investigates the use of quantum-inspired and fully quantum approaches for solving industry-relevant CFD problems. The studies validating tensor network methods across selected CFD systems provide strong evidence for the memory efficiency and long-term viability of this approach when applied to challenging flow problems (cf. Section 2.1). In this section we benchmarked three representative use-cases covering increasing physical and geometric complexity: flow around a cylinder, Rayleigh–Bénard convection, and flow through an S-bend geometry. These cases establish a systematic progression from laminar to turbulent dynamics and toward realistic engineering configurations, forming a rigorous validation baseline for later quantum extensions. To quantify the potential advantages of TNs, Section 2.1.1 introduces asymptotic advantage criteria, namely $\alpha < \beta/3$ for runtime advantage and $\alpha < \beta/2$ for memory advantage, where β characterizes the growth of grid resolution with Re and α comes the growth in complexity of the flow with increasing Re . Specifically, for the cylinder flow ($\beta \approx 0.8$), the measured exponents ($\alpha = 0.16\text{--}0.30$) already indicate clear memory gains and immediate potential for runtime benefits. In the case of Rayleigh–Bénard convection ($\beta \approx 1.2$), favorable memory scaling ($\alpha \approx 0.48\text{--}0.60$) is observed, even though a full runtime advantage isn't yet realized. However, previous work suggests a possible saturation of the bond dimension χ at higher Re , which, if confirmed here, would strongly support the long-term potential of MPS methods at larger Re , likewise for the quantum approach. Finally, for the S-bend geometry ($\beta \approx 1.7$), the compression results are exceptionally strong, revealing stable compression ($\chi < 10^2$ for a high fidelity of $\varepsilon_F < 10^{-10}$) and satisfying both asymptotic runtime and memory advantages with exponents $\alpha = 0.07\text{--}0.31$. Taken together, these results provide qualitative evidence that the MPS approach is a robust, memory-efficient alternative to classical solvers, thereby not only establishing a solid foundation for their integration into quantum frameworks but also allowing researchers to estimate classical lower bounds on potential quantum speedup.

Building on this, in Section 3.1 we motivate the transition from the MPS approach to a hybrid-classical approach via an VQA implementation. Namely in Section 3.1.1, we introduce the fundamental concepts of VQAs alongside a quantum version of Nyquist-Shannon theorem which provides a theoretical bound on the circuit width required to represent a function state. In Section 3.1.2, we demonstrated the potential of VQAs in representing 3D turbulent fields where we have shown that variational quantum circuits require less parameters than its MPS representations. Through scaling arguments, we show that we expect a computational advantage in the regime of large system size. Within the QCFD framework, MPOs are compiled into quantum circuits via tensor quantum programming (Section 3.1.3), directly connecting the classical and quantum representations. This mapping underpins the QNPU concept, which extends standard variational quantum architectures to handle nonlinear operations, such as those appearing in the Navier–Stokes equations, by combining multiple variational state copies with problem-specific linear operators. In this sense, MPS not only benchmarks and validates the quantum approach but also establish rigorous lower bounds on the achievable quantum advantage and the associated scaling behavior, which is expected to remain largely consistent in VQAs.

The prefactors of the scaling laws are largely neglected in this analysis and can only be accurately assessed through explicit VQA simulations and benchmarking on actual quantum hardware. Therefore, in Sections 3.1.4 and 3.1.5, we analyze VQAs and their resource scaling on two prototypical one-dimensional problems: the nonlinear Schrödinger equation and Burgers' equation. For the nonlinear Schrödinger equation (Section 3.1.4), the number of qubits scales as $n = 3 \log_2 N + 1$, while for Burgers' equation (Section 3.1.5), $n = 2 \log_2 N + 1$. In both cases, circuit depth and two-qubit gate counts are systematically optimized, with the Burgers' implementation achieving a threefold reduction in entangling gates compared to conventional designs. These findings confirm that nonlinear flow dynamics can be captured efficiently using shallow variational circuits, thereby improving the feasibility of practical quantum simulations. Section 3.1.6 discusses circuit-depth optimization strategies that trade additional qubits for reduced depth via mid-circuit measurements and classical control, yielding effectively non-unitary circuits that mitigate idling-qubit noise under current hardware constraints. Complementary to this, Section 3.1.7 proposes an approach to address barren plateaus (regions of exponentially vanishing gradients that hinder trainability) and demonstrates its potential on the one-dimensional quantum Heisenberg model.



Bibliography

- [1] Xiaojue Zhu et al. “Transition to the Ultimate Regime in Two-Dimensional Rayleigh–Bénard Convection”. In: *Phys. Rev. Lett.* 120 (14 Apr. 2018), p. 144502. DOI: 10.1103/PhysRevLett.120.144502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.144502>.
- [2] Tomohiro Hashizume et al. *Quantum computation at the edge of chaos*. 2026. eprint: arXiv: 2604.15441.
- [3] Nis-Luca van Hülst et al. *Quantum-Inspired Simulation of 2D Turbulent Rayleigh–Bénard Convection*. 2026. arXiv: 2604.16179 [physics.flu-dyn]. URL: <https://arxiv.org/abs/2604.16179>.
- [4] ENGYS Ltd. *HELIX: Open-source CFD software for enterprise*. version 4.4.0, <https://engys.com/helix/>. 2025. URL: <https://engys.com/helix/>.
- [5] Joel H. Ferziger, Milovan Perić, and Robert L. Street. *Computational Methods for Fluid Dynamics*. 2020. DOI: 10.1007/978-3-319-99693-6. URL: <https://app.dimensions.ai/details/publication/pub.1120368805>.
- [6] Hans Johnston and Charles R. Doering. “Comparison of Turbulent Thermal Convection between Conditions of Constant Temperature and Constant Flux”. In: *Phys. Rev. Lett.* 102 (6 Feb. 2009), p. 064501. DOI: 10.1103/PhysRevLett.102.064501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.102.064501>.
- [7] Nikita Gourianov et al. “A quantum-inspired approach to exploit turbulence structures”. In: *Nature Computational Science* 2.1 (Jan. 2022), pp. 30–37. ISSN: 2662-8457. DOI: 10.1038/s43588-021-00181-1. URL: <http://dx.doi.org/10.1038/s43588-021-00181-1>.
- [8] Leonhard Hölcher et al. “Quantum-inspired fluid simulation of two-dimensional turbulence with GPU acceleration”. In: *Physical Review Research* 7.1 (Jan. 2025). ISSN: 2643-1564. DOI: 10.1103/physrevresearch.7.013112. URL: <http://dx.doi.org/10.1103/PhysRevResearch.7.013112>.
- [9] Alexios A Michailidis, Christian Fenton, and Martin Kiffner. *Tensor Train Multiplication*. 2024. arXiv: 2410.19747 [physics.comp-ph]. URL: <https://arxiv.org/abs/2410.19747>.
- [10] Nis-Luca van Hülst et al. *Quantum-Inspired Tensor-Network Fractional-Step Method for Incompressible Flow in Curvilinear Coordinates*. 2025. arXiv: 2507.05222 [physics.flu-dyn]. URL: <https://arxiv.org/abs/2507.05222>.
- [11] Yuriel Núñez Fernández et al. “Learning tensor networks with tensor cross interpolation: New algorithms and libraries”. In: *SciPost Physics* 18.3 (Mar. 2025). ISSN: 2542-4653. DOI: 10.21468/scipostphys.18.3.104. URL: <http://dx.doi.org/10.21468/SciPostPhys.18.3.104>.
- [12] Hai-Yen Van. “Matrix Product State Simulations for Rayleigh–Bénard Convection”. Bachelor’s thesis. Hamburg, Germany: University of Hamburg, 2025.



- [13] Guenter Ahlers, Siegfried Grossmann, and Detlef Lohse. “Heat transfer and large scale dynamics in turbulent Rayleigh-Benard convection”. In: *Rev. Mod. Phys.* 81 (2009), pp. 503–537. DOI: 10.1103/RevModPhys.81.503. arXiv: 0811.0471 [physics.flu-dyn].
- [14] M. C. Cross and P. C. Hohenberg. “Pattern formation outside of equilibrium”. In: *Rev. Mod. Phys.* 65 (3 July 1993), pp. 851–1112. DOI: 10.1103/RevModPhys.65.851. URL: <https://link.aps.org/doi/10.1103/RevModPhys.65.851>.
- [15] Subrahmanyan Chandrasekhar. *Hydrodynamic and hydromagnetic stability*. Courier Corporation, 2013.
- [16] Chris A. Jones. “Planetary Magnetic Fields and Fluid Dynamos”. In: *Annual Review of Fluid Mechanics* 43.1 (Jan. 2011), pp. 583–614. DOI: 10.1146/annurev-fluid-122109-160727.
- [17] Geoffrey Vallis. *Atmospheric and Oceanic Fluid Dynamics*. Jan. 2006, p. 745. DOI: 10.2277/0521849691.
- [18] T. Kuhlbrodt et al. “On the driving processes of the Atlantic meridional overturning circulation”. In: *Reviews of Geophysics* 45.2 (2007). DOI: <https://doi.org/10.1029/2004RG000166>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2004RG000166>.
- [19] Eberhard Bodenschatz, Werner Pesch, and Guenter Ahlers. “Recent Developments in Rayleigh-Bénard Convection”. In: *Annual Review of Fluid Mechanics* 32 (2000), pp. 709–778. URL: <https://api.semanticscholar.org/CorpusID:38312855>.
- [20] Y.A. Çengel et al. *Fundamentals of Thermal-fluid Sciences*. McGraw-Hill Education, 2016. ISBN: 9789814720953. URL: <https://books.google.de/books?id=CssUvgAACAAJ>.
- [21] Francis H. Harlow and J. Eddie Welch. “Numerical Calculation of Time-Dependent Viscous Incompressible Flow of Fluid with Free Surface”. In: *The Physics of Fluids* 8.12 (Dec. 1965), pp. 2182–2189. ISSN: 0031-9171. DOI: 10.1063/1.1761178. eprint: https://pubs.aip.org/aip/pfl/article-pdf/8/12/2182/12731401/2182_1_online.pdf. URL: <https://doi.org/10.1063/1.1761178>.
- [22] Rung Thomas, Over Paul, Bengoechea Sergio, De Villiers Eugene and Scandurra Leonardo. *Quantum Computational Fluid Dynamics - WP Core Benchmark CFD Set: Deliverable 1.1. — Double-Bent Pipe — fdr.uni-hamburg.de*. <https://www.fdr.uni-hamburg.de/record/16637>. [Accessed 29-09-2025]. 2024.
- [23] Michael Lubasch et al. “Variational quantum algorithms for nonlinear problems”. In: *Physical Review A* 101.1 (2020), p. 010301. DOI: <https://doi.org/10.1103/PhysRevA.101.010301>. URL: <https://journals.aps.org/pr/abstract/10.1103/PhysRevA.101.010301>.
- [24] Dieter Jaksch et al. “Variational quantum algorithms for computational fluid dynamics”. In: *AIAA journal* 61.5 (2023), pp. 1885–1894. DOI: <https://doi.org/10.2514/1.J062426>. URL: <https://arc.aiaa.org/doi/abs/10.2514/1.J062426>.
- [25] Aram Harrow and Saeed Mehraban. “Approximate Unitary t -Designs by Short Random Quantum Circuits Using Nearest-Neighbor and Long-Range Gates”. In: *Communications in Mathematical Physics* 401.2 (July 2023), pp. 1531–1626. ISSN: 0010-3616, 1432-0916. DOI: 10.1007/s00220-023-04675-z. arXiv: 1809.06957 [quant-ph].
- [26] Jielun Chen, EM Stoudenmire, and Steven R White. “Quantum Fourier transform has small entanglement”. In: *PRX Quantum* 4.4 (2023), p. 040318.
- [27] Eric Perlman et al. “Data exploration of turbulence simulations using a database cluster”. In: *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*. 2007, pp. 1–11.

- [28] Yi Li et al. “A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence”. In: *Journal of Turbulence* 9 (2008), N31.
- [29] Wan Mingping et al. *Forced Isotropic Turbulence Data Set (Extended)*. 2012. DOI: 10.7281/T1KK98XB. URL: <https://turbulence.idies.jhu.edu/datasets/homogeneousTurbulence/isotropic>.
- [30] Pia Siegl et al. *Tensor-Programmable Quantum Circuits for Solving Differential Equations*. 2025. arXiv: 2502.04425 [quant-ph]. URL: <https://arxiv.org/abs/2502.04425>.
- [31] Martín Larocca et al. “Barren Plateaus in Variational Quantum Computing”. In: *Nature Reviews Physics* 7.4 (Apr. 2025), pp. 174–189. ISSN: 2522-5820. DOI: 10.1038/s42254-025-00813-9.
- [32] Daan Camps et al. *Explicit Quantum Circuits for Block Encodings of Certain Sparse Matrices*. 2023. arXiv: 2203.10236 [quant-ph]. URL: <https://arxiv.org/abs/2203.10236>.
- [33] A Termanova et al. “Tensor Quantum Programming”. In: *New Journal of Physics* 26.12 (Dec. 2024), p. 123019. DOI: 10.1088/1367-2630/ad985b.
- [34] Govind P. Agrawal. “Nonlinear Fiber Optics”. In: *Nonlinear Science at the Dawn of the 21st Century*. Ed. by P. L. Christiansen, M. P. Sørensen, and A. C. Scott. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 195–211. ISBN: 978-3-540-46629-1. DOI: <https://doi.org/10.1016/C2009-0-21165-2>. URL: <https://doi.org/10.1016/C2009-0-21165-2>.
- [35] K Nakkeeran. “Bright and dark optical solitons in fiber media with higher-order effects”. In: *Chaos, Solitons & Fractals* 13.4 (2002), pp. 673–679. DOI: [https://doi.org/10.1016/S0960-0779\(00\)00278-2](https://doi.org/10.1016/S0960-0779(00)00278-2). URL: <https://www.sciencedirect.com/science/article/abs/pii/S0960077900002782>.
- [36] Houria Triki et al. “Self-similar optical solitons with continuous-wave background in a quadratic-cubic non-centrosymmetric waveguide”. In: *Optics Communications* 437 (2019), pp. 392–398. DOI: <https://doi.org/10.1016/j.optcom.2018.12.074>. URL: <https://www.sciencedirect.com/science/article/abs/pii/S0030401818311295>.
- [37] Catherine Sulem and Pierre-Louis Sulem. “The nonlinear Schrödinger equation: self-focusing and wave collapse”. In: vol. 139. Springer Science & Business Media, 1999. DOI: <https://doi.org/10.1007/b98958>. URL: <https://doi.org/10.1007/b98958>.
- [38] Eugene P Gross. “Structure of a quantized vortex in boson systems”. In: *Il Nuovo Cimento (1955-1965)* 20.3 (1961), pp. 454–477. DOI: <https://doi.org/10.1007/BF02731494>. URL: <https://link.springer.com/article/10.1007/BF02731494>.
- [39] Lev P Pitaevskii. “Vortex lines in an imperfect Bose gas”. In: *Sov. Phys. JETP* 13.2 (1961), pp. 451–454. URL: http://www.jetp.ras.ru/cgi-bin/dn/e_013_02_0451.pdf.
- [40] Franco Dalfovo et al. “Theory of Bose-Einstein condensation in trapped gases”. In: *Reviews of modern physics* 71.3 (1999), p. 463. DOI: <https://doi.org/10.1103/RevModPhys.71.463>. URL: <https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.71.463>.
- [41] Anthony J Leggett. “Bose-Einstein condensation in the alkali gases: Some fundamental concepts”. In: *Reviews of modern physics* 73.2 (2001), p. 307. DOI: <https://doi.org/10.1103/RevModPhys.73.307>. URL: <https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.73.307>.
- [42] L. Pitaevskii and S. Stringari. *Bose-Einstein Condensation*. 2003.
- [43] Mark Edwards and Keith Burnett. “Numerical solution of the nonlinear Schrödinger equation for small samples of trapped neutral atoms”. In: *Physical Review A* 51.2 (1995), p. 1382. DOI: <https://doi.org/10.1103/PhysRevA.51.1382>. URL: <https://journals.aps.org/pra/abstract/10.1103/PhysRevA.51.1382>.



- [44] Michael Lubasch, Pierre Moinier, and Dieter Jaksch. “Multigrid renormalization”. In: *Journal of Computational Physics* 372 (2018), pp. 587–602. DOI: <https://doi.org/10.1016/j.jcp.2018.06.065>. URL: <https://www.sciencedirect.com/science/article/abs/pii/S0021999118304431>.
- [45] Muhammad Umer et al. “Probing the limits of variational quantum algorithms for nonlinear ground states on real quantum hardware: The effects of noise”. In: *Phys. Rev. A* 111 (1 Jan. 2025), p. 012626. DOI: [10.1103/PhysRevA.111.012626](https://doi.org/10.1103/PhysRevA.111.012626). URL: <https://link.aps.org/doi/10.1103/PhysRevA.111.012626>.
- [46] Muhammad Umer, Eleftherios Mastorakis, and Dimitris G Angelakis. “Efficient estimation and sequential optimization of cost functions in variational quantum algorithms”. In: *Quantum Science and Technology* 10.3 (May 2025), p. 035022. DOI: [10.1088/2058-9565/add55e](https://doi.org/10.1088/2058-9565/add55e). URL: <https://dx.doi.org/10.1088/2058-9565/add55e>.
- [47] Jérémie Bec and Konstantin Khanin. “Burgers turbulence”. In: *Physics reports* 447.1-2 (2007), pp. 1–66. DOI: <https://doi.org/10.1016/j.physrep.2007.04.002>. URL: <https://doi.org/10.1016/j.physrep.2007.04.002>.
- [48] Eleftherios Mastorakis et al. “Resource-Efficient Hadamard Test Circuits for Nonlinear Dynamics on a Trapped-Ion Quantum Computer”. In: *arXiv preprint arXiv:2507.19250* (2025). DOI: <https://doi.org/10.48550/arXiv.2507.19250>. URL: <https://doi.org/10.48550/arXiv.2507.19250>.

