

Simulating sparse SYK model with a randomized algorithm on a trapped-ion quantum computer

Received: 17 July 2025

Accepted: 12 February 2026

Cite this article as: Granet, E., Kikuchi, Y., Dreyer, H. *et al.* Simulating sparse SYK model with a randomized algorithm on a trapped-ion quantum computer. *npj Quantum Inf* (2026). <https://doi.org/10.1038/s41534-026-01206-1>

Etienne Granet, Yuta Kikuchi, Henrik Dreyer & Enrico Rinaldi

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

Simulating sparse SYK model with a randomized algorithm on a trapped-ion quantum computer

Etienne Granet,^{1,*} Yuta Kikuchi,^{2,3} Henrik Dreyer,¹ and Enrico Rinaldi^{4,5}

¹*Quantinuum, Leopoldstrasse 180, 80804 Munich, Germany*

²*Quantinuum K.K., Otemachi Financial City Grand Cube 3F, 1-9-2 Otemachi, Chiyoda-ku, Tokyo, Japan*

³*RIKEN Center for Interdisciplinary Theoretical and Mathematical Sciences (iTHEMS), RIKEN, Wako, Saitama 351-0198, Japan*

⁴*Quantinuum, Partnership House, Carlisle Place, London SW1P 1BX, UK*

⁵*RIKEN Center for Quantum Computing (RQC), RIKEN, Wako, Saitama 351-0198, Japan*

(Dated: February 12, 2026)

The Sachdev-Ye-Kitaev (SYK) model describes a strongly correlated quantum system that shows a strong signature of quantum chaos. Due to its chaotic nature, the simulation of real-time dynamics becomes quickly intractable by means of classical numerics, and thus, quantum simulation is deemed to be an attractive alternative. Nevertheless, quantum simulations of the SYK model on noisy quantum processors are severely limited by the complexity of its Hamiltonian. In this work, we simulate the real-time dynamics of a sparsified version of the SYK model with 24 Majorana fermions on a trapped-ion quantum processor. We adopt a randomized quantum algorithm, TETRIS, and develop an error mitigation technique tailored to the algorithm. Leveraging the hardware’s high-fidelity quantum operations and all-to-all connectivity of the qubits, we successfully calculate the Loschmidt amplitude for sufficiently long times so that its decay is observed. Based on the experimental and further numerical results, we assess the future possibility of larger-scale simulations of the SYK model by estimating the required quantum resources. Moreover, we present a scalable mirror-circuit benchmark based on the randomized SYK Hamiltonian and the TETRIS algorithm, which we argue provides a better estimate of the decay of fidelity for local observables than standard mirror-circuits.

INTRODUCTION

The Sachdev-Ye-Kitaev (SYK) model describes a strongly-interacting quantum system that consists of N Majorana fermions with random q -body interactions [1–3]. Its connections to condensed matter physics and non-Fermi liquids make it an attractive model to study strong electronic correlations and disorder in metals and cuprates [4–6]. The model has also gained significant attention as it admits a two-dimensional gravitational holographic description in the infrared regime. In fact, it has been used as a paradigmatic model for “Quantum Gravity in the Lab” [7–11], a research effort to study quantum gravity phenomena through the holographic duality lenses, leading to new understandings (see e.g. [6, 12, 13] for reviews). The SYK model exhibits a strong signature of quantum chaos by saturating a universal bound of quantum Lyapunov exponent at low temperatures [14]. Despite many analytical studies that address the static and dynamic properties of the SYK model under certain limits, further numerical investigations are required to understand them for the other regimes, e.g., finite N , finite-range interaction, etc. In particular, simulation of real-time dynamics is ubiquitous in addressing signatures of quantum chaos through the Loschmidt echo, out-of-time-ordered correlators, spectral form factors, and other observables in the toolbox of condensed matter physics and quantum information science.

In this work, we simulate the real-time dynamics of the SYK model with four-body interactions ($q = 4$) to

explore its feasibility on noisy quantum hardware. We remark that digital quantum simulations of the SYK models on noisy quantum hardware have been severely limited due to their fully non-local interactions and the number of terms scaling as $O(N^4)$ (see [4, 15, 16] for experimental works). We alleviate these issues by considering a sparsified version of the SYK model [17–23] and employing a recently developed randomized Hamiltonian simulation algorithm, TETRIS (Time Evolution Through Random Independent Sampling) [24]. This stochastic algorithm enables Hamiltonian simulation without discretization error and is particularly suited for simulating the SYK model, leveraging the randomized nature of the SYK Hamiltonian.

As a demonstrative example we focus on a single dynamical observable: the vacuum survival probability, or Loschmidt amplitude. We choose the Loschmidt amplitude [25–28] because it is a key observable to study quantum chaos in the SYK model due to its known connection [27, 28] with out-of-time-ordered correlators (OTOC) [29–34]. In our experiment, we evolve the system for time t under the sparse SYK dynamics employing TETRIS and calculate the survival probability of the all-zero state on a trapped-ion quantum processor, whose all-to-all qubit connectivity facilitates the simulation of non-local interaction in the model. We carefully diagnose the error sources contributing to the measurement outcomes. Moreover, we develop an error mitigation technique that relies on properties of the randomized algorithm we use for computing the time evolution. Our protocol enables

the successful simulation of $N = 24$ sparse SYK model until the probability decay inherent to the model's ideal dynamics is observed at $Jt \sim 1$. We remark that the decay of the Loschmidt amplitude with time is not necessarily specific to chaotic models. For instance, even integrable models can decay at early times. However, chaotic models such as SYK models do not exhibit amplitude revival in times polynomial in the system size.

Besides simulating quantum dynamics of the SYK model, we propose to employ the TETRIS algorithm for benchmarking the noise effect of a quantum processor on a circuit used to estimate a local observable. Mirror circuit benchmarks provide scalable benchmarking protocols, where one measures the survival probability of the initial state after an application of a mirror circuit, that is, a unitary circuit followed by its inverse [35–37]. The performance of a quantum processor in running a circuit with a given number of gates is measured by the decay of probability from 1, the ideal value, with faster decays indicating stronger noise effects. In the present work, we consider a variant of the mirror circuit benchmarks, where we use two independently sampled TETRIS unitary circuits, U and U' , that agree only on average with the correct time evolution operator. Thus, the circuit $U'^{\dagger}U$ provides a mirror-on-average circuit in the sense that it is proportional to the identity operator only on average [38]. We argue that this benchmark gives a better estimate of noise on local observables than the standard mirror circuit benchmark when using this stochastic algorithm for implementing the time evolution.

Finally, we conclude by estimating the required quantum resources by extrapolating our numerical and experimental results and address the feasibility of larger simulations of the SYK model.

RESULTS

Setup

The $q = 4$ SYK model is a quantum mechanical model with N Majorana fermions ψ_i ($\{\psi_i, \psi_j\} = 2\delta_{ij}$) described by the Hamiltonian [2, 3],

$$H = \sum_{i < j < k < l} J_{ijkl} \psi_i \psi_j \psi_k \psi_l, \quad (1)$$

The couplings J_{ijkl} are independent random Gaussian variables with mean 0 and variance $\text{Var}[J_{ijkl}] = \frac{3!J^2}{N^3}$. The sparse SYK model is defined by [17–23]

$$H = \sum_{i < j < k < l} p_{ijkl} J_{ijkl} \psi_i \psi_j \psi_k \psi_l, \quad (2)$$

where $p_{ijkl} \in \{0, 1\}$ are randomly sampled, being equal to 1 with probability p , where p controls the sparsity of

the model. The average number of terms in the Hamiltonian (2) is $p \binom{N}{4} = O(pN^4)$. For gravitational physics to emerge at the infrared regime, one needs to retain an extensive number of terms on average by setting the probability to $p = kN / \binom{N}{4}$ for an N -independent sparsity parameter k [17]. To maintain the extensive energy, the variance is rescaled as $\text{Var}[J_{ijkl}] = \frac{3!J^2}{pN^3}$. Upon taking the disorder average, we find the 1-norm of the Hamiltonian $\|H\|_1 = \sum_{i < j < k < l} p_{ijkl} |J_{ijkl}|$ is

$$\frac{\sqrt{6p}J}{24} \frac{N!}{N^{3/2}(N-4)!} = O(\sqrt{p}N^{5/2}) = O(N). \quad (3)$$

This is obtained from the fact that the mean of p_{ijkl} is p , the mean of $|J_{ijkl}|$ is of order $\sqrt{\text{Var}[J_{ijkl}]}$, and the sum contains $\binom{N}{4}$ terms. This Hamiltonian on N Majorana fermions can be expressed in terms of Pauli matrices through the Jordan-Wigner transformation into $L = N/2$ qubits

$$\begin{aligned} \psi_{2j} &= X_j Z_{j-1} Z_{j-2} \dots Z_1, \\ \psi_{2j+1} &= Y_j Z_{j-1} Z_{j-2} \dots Z_1. \end{aligned} \quad (4)$$

We write the corresponding spin Hamiltonian as

$$H = \sum_n c_n P_n, \quad (5)$$

with a Pauli string P_n . The resulting SYK Hamiltonians are thus all-to-all coupled, with long Pauli strings of average length $O(L)$, and with random coefficients that can take both small and large values. These properties make these Hamiltonians difficult to simulate. Note that a different fermion encoding could be used, but local encodings [39, 40] provide diminishing returns due to the non-local and all-to-all nature of the interactions. The ternary tree encoding [41] could be an attractive possibility due to its favourable scaling, despite the need for additional ancillas, and we will explore it later on. For the system size studied $N = 24$, we found that the Jordan-Wigner encoding was more efficient.

The most widely used simulation technique, Trotterization, scales particularly badly with L for these systems. It could be rather complicated to derive tight bounds on the gate complexity and Trotter error. For a recent work on higher-order product formulas focusing on the SYK model, see Ref. [42]. The time evolution operator is approximated by s Trotter steps using the first-order Trotter decomposition,

$$e^{iHt} \approx \left(e^{ic_1 P_1 t/s} \dots e^{ic_M P_M t/s} \right)^s. \quad (6)$$

Since the Pauli strings are of length $O(L)$ on average and the exponentiation of a Pauli string of length $O(L)$ requires $O(L)$ two-qubit (TQ) gates, the cost to implement one single Trotter step is $O(pL^5)$ TQ gates. Equation (6) comes with an discretization (Trotter) error $O(t^2 pL^4/s)$.

Indeed, the Trotter error arises from non-vanishing commutators of the form $[\psi_i\psi_j\psi_k\psi_l, \psi_{i'}\psi_{j'}\psi_{k'}\psi_{l'}]$. For this to be finite, two four-body terms must share at least one of the fermion indices. There are $O(N^7)$ such combinations of $(ijkl)$ and $(i'j'k'l')$. Moreover, each four-body term labelled by $(ijkl)$ exists in a single disorder realization of the Hamiltonian (2) with probability p . Therefore, the number of non-vanishing commutators scales as $O(p^2N^7)$, resulting in the Trotter error $t^2/s \times \text{Var}[J_{ijkl}] \times O(p^2N^7) = O(t^2pN^4/s)$. Thus, the total TQ gate cost is $O(p^2L^9t^2/\epsilon)$ to achieve the Trotter error ϵ . Setting $p \sim L^{-3}$, as required to reproduce the properties of the dense SYK model, we find a gate count $O(L^3t^2/\epsilon)$.

Various randomized algorithms for Hamiltonian simulation provide favourable quantum resources compared to the Trotterization in a number of cases [43–50]. As we will see below, the randomized TETRIS algorithm [24] involves $t^2\mu^2$ rotations of Pauli strings, where μ is the 1-norm of the Hamiltonian. Most notably, the gate cost of each circuit is independent of the required error ϵ . Since each rotation requires $O(L)$ TQ gates, this yields a TQ gate count $O(\mu^2t^2L) = O(pL^6t^2)$. The scaling is strictly better than the Trotterization scaling for any sparsity parameter $p \sim L^{-\alpha}$ with $\alpha < 3$. For $p \sim L^{-3}$, TETRIS achieves no discretization error at the same cost as the Trotterization. Furthermore, no matter how small the evolution time t is, the Trotterization needs to apply at least one Trotter step, which costs $O(pL^5)$. On the other hand, the TETRIS gate cost decreases as t^2 , and thus, becomes cheaper for short-time evolution. This will be numerically verified later in Fig. 6. To give some precise numbers beyond this scaling analysis, for $N = 24$ Majorana fermions, we get an average number of TQ gates around 1022 per Trotter step, before circuit optimization. For the TETRIS algorithm at optimal gate angle, we get an average number of TQ gates around $3100t^2$ for simulation time t before circuit optimization.

The TETRIS algorithm takes as a parameter a gate angle $0 < \tau < \pi/2$, and generates random unitary operators U that average to

$$\mathbb{E}_U[U] = \lambda e^{iHt}, \quad (7)$$

where \mathbb{E}_U denotes the statistical average over random samples of U , and λ is

$$\lambda = e^{-t\mu \tan(\tau/2)}, \quad (8)$$

with the 1-norm $\mu = \sum_n |c_n|$ of the Hamiltonian (5). The random unitary operator U is expressed as a product $U = e^{i\tau P_{j_1}} \dots e^{i\tau P_{j_M}}$, where the multiplication by $e^{i\tau P_n}$ are random events that follow independent Poisson processes with rates $|c_n|/\sin \tau$ during an evolution time t . Namely, the unitary operator U is built starting from the identity $U = I$, and for each time window of small duration dt between time 0 and time t , we multiply U

by $e^{i\tau P_n}$ with a probability $dt|c_n|/\sin \tau$. We refer to [24] for a detailed description of the algorithm. As a consequence, the average number of rotation gates per sampled unitary operator U is $t\mu/\sin \tau$. The attenuation factor λ increases the number of shots to reach a given precision by a factor $1/\lambda^2 = e^{2t\mu \tan(\tau/2)}$. The optimal choice of τ that minimizes the total number of gates across different shots is $\tau \approx 1/(t\mu)$ for large $t\mu$. This yields an average number of rotations $t^2\mu^2$. Note that for the SYK Hamiltonian, the long Pauli strings that will be exponentiated result mostly in TQ gates with maximal angle, coming from the decomposition of Pauli gadgets.

Noise mitigation: Echo verification

The implementation of TETRIS requires an ancillary qubit to calculate the average of unitary operators U . We discuss an error mitigation scheme that makes use of the measurement outcomes on the system register [51–54]. Let us consider a density matrix ρ on the ancilla and the system qubits, that we initialize in $|+\rangle\langle+|$ for the ancilla, and $|0\rangle\langle 0| := |0\dots 0\rangle\langle 0\dots 0|$ for the system qubits. After application of a unitary operator U conditioned on the ancilla being $|1\rangle$, we have

$$\begin{aligned} \rho &= \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| \\ &+ \frac{1}{2}|1\rangle\langle 0| \otimes U|0\rangle\langle 0| + \frac{1}{2}|0\rangle\langle 1| \otimes |0\rangle\langle 0|U^\dagger \\ &+ \frac{1}{2}|1\rangle\langle 1| \otimes U|0\rangle\langle 0|U^\dagger, \end{aligned} \quad (9)$$

where the registers before and after \otimes denote the ancilla and the system qubits respectively. Averaging over U , we get

$$\begin{aligned} \rho_{\text{ave}} &:= \mathbb{E}_U[\rho] \\ &= \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| \\ &+ \frac{\lambda}{2}|1\rangle\langle 0| \otimes e^{iHt}|0\rangle\langle 0| + \frac{\lambda}{2}|0\rangle\langle 1| \otimes |0\rangle\langle 0|e^{-iHt} \\ &+ \frac{1}{2}|1\rangle\langle 1| \otimes \mathbb{E}_U[U|0\rangle\langle 0|U^\dagger]. \end{aligned} \quad (10)$$

One notes that the average of $U|0\rangle\langle 0|U^\dagger$ is a priori unknown since it is quadratic in U . Measuring X on the ancilla and \mathcal{O} , an arbitrary observable, on the system qubits, we have

$$\text{Tr}[\rho_{\text{ave}}(X \otimes \mathcal{O})] = \lambda \Re[\langle 0|\mathcal{O}e^{iHt}|0\rangle]. \quad (11)$$

The real part of the Loschmidt amplitude can thus be obtained by setting $\mathcal{O} = I$, i.e., only measuring the ancilla

$$\Re[\langle 0|e^{iHt}|0\rangle] = \lambda^{-1}\text{Tr}[\rho_{\text{ave}}(X \otimes I)]. \quad (12)$$

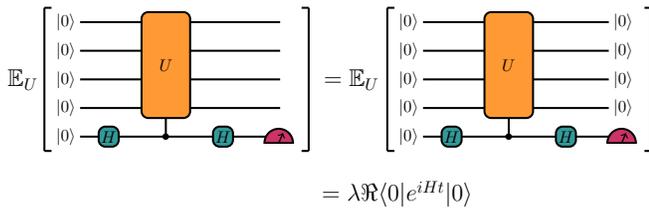


FIG. 1. Schematic circuit representation for the computation of the Loschmidt echo amplitude and the echo verification.

However, it can also be equivalently obtained by setting $\mathcal{O} = |0\rangle\langle 0|$, the projector onto the initial state,

$$\Re[\langle 0|e^{iHt}|0\rangle] = \lambda^{-1}\text{Tr}[\rho_{\text{ave}}(X \otimes |0\rangle\langle 0|)]. \quad (13)$$

When measuring the system qubits in the Z basis, the difference between the two cases above is that when computing $X \otimes I$, the measurement outcome of the system qubits is not taken into account, whereas when computing $X \otimes |0\rangle\langle 0|$, we count zero whenever the measurement outcome of the system qubits is not 0...0. This is different from discarding the corresponding shots because the overall number of shots remains the same. Similarly, projecting the system qubits onto any product state in the Z basis that is different from $|0\dots 0\rangle$, i.e., taking $\mathcal{O} = |i_1\dots i_L\rangle\langle i_1\dots i_L|$ has expectation value 0 whenever $i_1\dots i_L \neq 0\dots 0$. It follows that counting the shots where the system qubits are not 0...0 does not modify the expectation values, but increases the shot noise. On a noiseless quantum computer, it is thus always more efficient to measure $X \otimes |0\rangle\langle 0|$. On a noisy quantum computer, measuring $X \otimes |0\rangle\langle 0|$ allows for discarding some bit flips on the system qubits, at the cost of reducing the number of shots available and thus increasing sampling overhead. This noise mitigation technique is called echo verification [51, 54]. It is represented schematically in Fig 1. In panel (a) of Fig. 2, we show the variance over circuits, for different numbers of shots per circuit and different numbers of gates, and for $\mathcal{O} = I$, $\mathcal{O} = |0\rangle\langle 0|$. We observe that $\mathcal{O} = |0\rangle\langle 0|$ has generally smaller variance than $\mathcal{O} = I$, as expected. As for the influence of the number of shots per circuit, it is known that doing 1 shot per random TETRIS circuit minimizes the variance at fixed total number of shots, when neglecting compilation cost [24]. We observe that in practice, doing 10 shots per circuit reduces the standard deviation by around $\sqrt{10}$, which is thus close to optimal. On the contrary, doing 10^4 shots per random circuit reduces the standard deviation by much less than $\sqrt{10^4}$ which is far from the optimal case.

At early times before the Loschmidt amplitude decays, $|0\dots 0\rangle$ has the dominant amplitude in an evolving state. On a noisy quantum computer, some bit flips can occur to change the state $|0\dots 0\rangle$ to a state of a small Hamming weight. Counting 0 for these states will thus attenuate the signal. Since fermionic parity is conserved in the

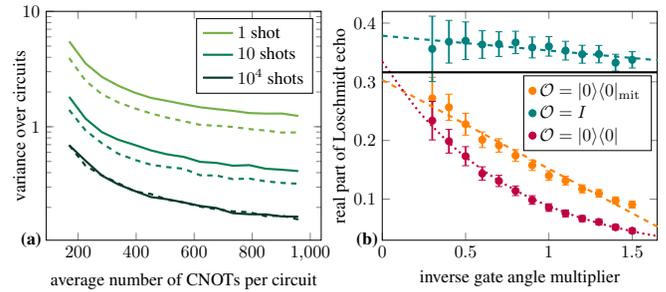


FIG. 2. **a:** Standard deviation of measurements of real part of the Loschmidt amplitude at $Jt = 0.5$, as a function of number of CNOTs per circuit (without any circuit optimization), for different number of shots per circuits (different shades of green) and measurement operator $\mathcal{O} = I$ (solid lines) and $\mathcal{O} = |0\rangle\langle 0|$ (dashed lines). The circuits are sampled both from TETRIS and from different disorder realizations of the SYK Hamiltonian. **b:** Real part of Loschmidt amplitude as a function of $1/\alpha$ for gate angle $\tau = \alpha\tau_0$ with $\tau_0 = 1/(t\mu)$, for $Jt = 0.5$, simulated with a depolarizing noise channel of amplitude 0.001 after every TQ gate, for different operators \mathcal{O} . The error bars correspond to 10^3 circuits with 10 shots each. The dashed lines are linear fits and the dotted line is an exponential fit. The black line shows the exact value.

SYK model, shots with an odd number of bit flips can be attributed to noise with certainty. For a state with a small Hamming weight, noise is more likely to increase the weight than to decrease it. Therefore, the majority of the shots with a single 1 are expected to come from a bit-flip on top of 0...0 than on top of a bitstring with Hamming weight 2. Taking these shots into account should thus improve the estimate of $X \otimes |0\rangle\langle 0|$ on a noisy quantum computer. To account for this, we propose a variant of the echo verification by measuring the following observable

$$|0\rangle\langle 0|_{\text{mit}} := |0\dots 0\rangle\langle 0\dots 0| + \sum_{j=1}^L |0\dots 010\dots 0\rangle\langle 0\dots 010\dots 0|, \quad (14)$$

where the single 1 is at position j in each bitstring.

Hardware results

We now present a hardware implementation of our setup on Quantinuum System Model H1 [55]. This is a trapped-ion quantum computing device hosting 20 all-to-all coupled qubits. We fix the number of Majorana fermions to $N = 24$, corresponding to $L = 12$ qubits, and add one ancilla qubit, using thus 13 qubits in total. The sparsity parameter is $k = 2.3$, below which the ramp, a universal behaviour predicted by random matrix theory, disappears in the spectral form factor, indicating the loss of spectral rigidity [23]. Therefore, with $k = 2.3$ we choose the most sparse SYK that could still be com-

time	shallow	deep
$t = 0.2$	940	140
$t = 0.35$	820	105
$t = 0.5$	440	50
$t = 0.65$	285	30
$t = 0.8$	180	21

TABLE I. Number of random circuits generated and run on hardware, for different times, with 6 shots performed per random circuit.

patible with chaotic properties related to its holographic dual gravity model. We do not intend this experiment to be a probe of a realistic gravitational dual model, but to provide a lower bound on the number of gates needed to simulate a SYK model with chaotic properties. On the hardware, we run 6 shots per circuit, each circuit being randomly sampled among different TETRIS unitary operators U and among different disorder realizations of the sparse SYK Hamiltonian H (2). We use the ability of the hardware to perform mid-circuit measurement and qubit reset to “stitch” multiple circuits in a same run, in order to minimize the overhead cost in running several circuits with few shots each. We consider two different gate angles, $\tau_0 = 1.5/(t\mu)$, where μ is the 1-norm of the sampled Hamiltonian H , that we refer to as “shallow circuits” and one equal to $\tau_\alpha = \alpha\tau_0$ with $\alpha = 1/3$, that we refer to as “deep circuits”. The deep circuits contain, on average, around three times more gates than the shallow circuits. Taking a large difference of number of gates in the two circuits used to perform ZNE reduces the shot noise amplification by the extrapolation [56]. As described in the Methods section, this allows us to apply Large Gate Angle Extrapolation (LGAE) on the shallow and deep circuits and to study the efficiency of the three different measurement strategies, where $\mathcal{O} = I, |0\rangle\langle 0|, |0\rangle\langle 0|_{\text{mit}}$. We simulate the system up to time $Jt = 0.8$, where the average real part of the Loschmidt echo starts to be close to 0, as seen in panel (b) of Fig 4. In total, the number of circuits run on hardware are given in Table I. All hardware results are summarized in Fig. 3.

We observe that for $\mathcal{O} = |0\rangle\langle 0|_{\text{mit}}$, the LGAE mitigation recovers the exact values within just about one standard deviation. In contrast, for $\mathcal{O} = |0\rangle\langle 0|$ the mitigation leads to an underestimation of the exact result. We attribute it to the fact that, because this observable is noisier, we are beyond the regime where noise effects are linear in the number of gates, as we already saw in panel (b) of Fig. 2.

In Fig. 4 we plot the mitigated results using the exponential form of LGAE. As expected, the error bars are significantly larger. Here, even the case $\mathcal{O} = |0\rangle\langle 0|$ correctly recovers the exact result within error bars, although the error bars are too large to be really conclusive for $Jt = 0.65$ and $Jt = 0.8$. For $\mathcal{O} = I$ with the linear extrapolation in Fig. 3, the results display

significantly large error bars, with both shallow and deep circuits giving comparable results. In that case the efficiency of the noise mitigation is also less conclusive due again to the large error bars. We also observe that the signal is typically *larger* than the exact values. This *upward* bias due to noise is not standard because noise typically biases the results towards vanishing signal, except for non-unital channels such as leakage error. We next study theoretically the origin of this effect.

Effect of noise

We now discuss the effect of noise on the expectation value $\text{Tr}[\rho_{\text{ave}}(X \otimes \mathcal{O})]$. We use $\langle X \otimes \mathcal{O} \rangle_{\text{noisy}}$ to denote the noisy expectation values, while $\text{Tr}[\rho_{\text{ave}}(X \otimes \mathcal{O})]$ denotes the noiseless expectation value.

We analyze the dominant source of error occurring on the system qubits. To get tractable expressions, we model the noise effects on the system qubits with a global depolarizing channel \mathcal{N} that occurs with rate q . If the channel \mathcal{N} is applied at time t_1 , the density matrix (10) is turned into

$$\begin{aligned} \mathcal{N}[\rho_{\text{ave}}] = & \\ & \frac{1}{2}|0\rangle\langle 0| \otimes \frac{I^N}{2^N} + \frac{\lambda_1}{2}|1\rangle\langle 0| \otimes \frac{I^N}{2^N} \langle 0|e^{iHt_1}|0\rangle \\ & + \frac{\lambda_1}{2}|0\rangle\langle 1| \otimes \frac{I^N}{2^N} \langle 0|e^{-iHt_1}|0\rangle + \frac{1}{2}|1\rangle\langle 1| \otimes \frac{I^N}{2^N}, \end{aligned} \quad (15)$$

with $\lambda_1 = e^{-t_1\mu \tan(\tau/2)}$ the attenuation factor up to time t_1 . Further evolving the state until time t_2 , we find

$$\begin{aligned} & \frac{1}{2}|0\rangle\langle 0| \otimes \frac{I^N}{2^N} + \frac{\lambda_2}{2}|1\rangle\langle 0| \otimes \frac{e^{iH(t_2-t_1)}}{2^N} \langle 0|e^{iHt_1}|0\rangle \\ & + \frac{\lambda_2}{2}|0\rangle\langle 1| \otimes \frac{e^{-iH(t_2-t_1)}}{2^N} \langle 0|e^{-iHt_1}|0\rangle + \frac{1}{2}|1\rangle\langle 1| \otimes \frac{I^N}{2^N}, \end{aligned} \quad (16)$$

by averaging over the TETRIS circuits, with $\lambda_2 = e^{-t_2\mu \tan(\tau/2)}$. Hence, if n errors occurs at times $0 < t_1 < \dots < t_n < t$, we measure the contribution at time t

$$\begin{aligned} & \lambda \langle 0|e^{iHt_1}|0\rangle \frac{\text{Tr}[e^{iH(t_2-t_1)}]}{2^N} \times \\ & \frac{\text{Tr}[e^{iH(t_3-t_2)}]}{2^N} \dots \frac{\text{Tr}[e^{iH(t_n-t_{n-1})}]}{2^N} \frac{\text{Tr}[\mathcal{O}e^{iH(t-t_n)}]}{2^N}. \end{aligned} \quad (17)$$

Since each error occurs randomly and independently from other errors, the statistics of number of errors within a time window follows a Poisson distribution. So the probability of having n errors during a time t is $\frac{(qt)^n}{n!} e^{-qt}$. The measured value constrained to having n errors is expressed by the integral of (17) over $0 < t_1 < \dots < t_n < t$ with a normalization factor $n!/t^n$. Thus for example,

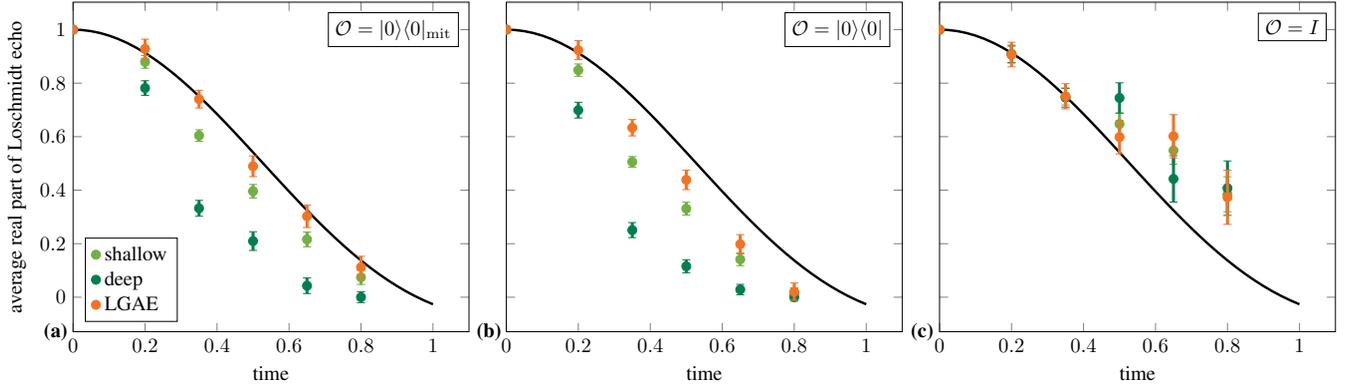


FIG. 3. Hardware results from Quantinuum H1-1 ion-trap quantum computer. Real part of the Loschmidt amplitude, averaged over random realizations of the SYK Hamiltonian for $N = 24$ fermions, as a function of time Jt , for shallow and deep circuits, together with the LGAE mitigation technique, for the three tested measurement operators \mathcal{O} . The black line shows the exact value. The errors bars correspond to one standard deviation computed across different shots and random circuit realizations.

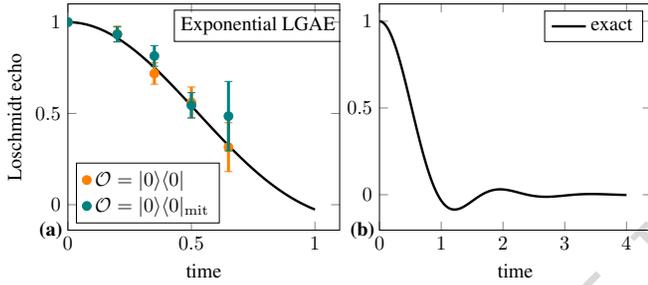


FIG. 4. **a**: Hardware results after exponential LGAE, in the same setting as in Fig. 3. The mitigated results for $Jt = 0.8$ have very large variance and are outside of the plotted window. **b**: Exact average real part of the Loschmidt echo as a function of time, for longer times, showing no revivals.

at order $(qt)^2$ we have

$$\begin{aligned}
 \langle X \otimes \mathcal{O} \rangle_{\text{noisy}} + i \langle Y \otimes \mathcal{O} \rangle_{\text{noisy}} &= \lambda e^{-qt} \langle 0 | \mathcal{O} e^{iHt} | 0 \rangle \\
 &+ \lambda q e^{-qt} \int_0^t dt_1 \langle 0 | e^{iHt_1} | 0 \rangle \frac{\text{Tr}[\mathcal{O} e^{iH(t-t_1)}]}{2^N} \\
 &+ \lambda q^2 e^{-qt} \int_0^t dt_1 \int_0^{t_1} dt_2 \langle 0 | e^{iHt_2} | 0 \rangle \times \\
 &\quad \frac{\text{Tr}[e^{iH(t_1-t_2)}]}{2^N} \frac{\text{Tr}[\mathcal{O} e^{iH(t-t_1)}]}{2^N} \\
 &+ O((qt)^3).
 \end{aligned} \tag{18}$$

The first line corresponds to the contribution of 0 errors, the second line to 1 error occurring at time t_1 , the third line to 2 errors occurring at times t_1 and t_2 . Collecting the contributions with arbitrary numbers of errors, it can be concisely written as

$$\begin{aligned}
 \langle X \otimes \mathcal{O} \rangle_{\text{noisy}} + i \langle Y \otimes \mathcal{O} \rangle_{\text{noisy}} \\
 = e^{-qt} \lambda \langle 0 | \mathcal{O} e^{iHt} | 0 \rangle + \lambda q e^{-qt} \int_0^t dt_1 \frac{\text{Tr}[\mathcal{O} e^{iH(t-t_1)}]}{2^N} F(t_1),
 \end{aligned} \tag{19}$$

by introducing $F(t)$ satisfying the integral equation

$$F(t) = \langle 0 | e^{iHt} | 0 \rangle + q \int_0^t ds \frac{\text{Tr}[e^{iH(t-s)}]}{2^N} F(s). \tag{20}$$

Indeed, inserting repeatedly n times $F(t)$ in (19) yields the contribution to up to n errors.

At a low error rate q and for $\mathcal{O} = |0\rangle\langle 0|$, we have

$$\begin{aligned}
 \langle X \otimes |0\rangle\langle 0| \rangle_{\text{noisy}} &= \lambda \Re \langle 0 | e^{iHt} | 0 \rangle \\
 &+ \lambda q \Re \left(\int_0^t dt_1 \frac{\langle 0 | e^{iHt_1} | 0 \rangle \langle 0 | e^{iH(t-t_1)} | 0 \rangle}{2^N} - t \langle 0 | e^{iHt} | 0 \rangle \right) \\
 &+ O((qt)^2).
 \end{aligned} \tag{21}$$

Because of the factor $1/2^N$, the parenthesis is very likely to be negative at small t . Hence the signal is decreased. On the other hand, when $\mathcal{O} = I$, we have

$$\begin{aligned}
 \langle X \otimes I \rangle_{\text{noisy}} &= \lambda \Re \langle 0 | e^{iHt} | 0 \rangle \\
 &+ \lambda q \Re \left(\int_0^t dt_1 \frac{\text{Tr}[e^{iHt_1}]}{2^N} \langle 0 | e^{iH(t-t_1)} | 0 \rangle - t \langle 0 | e^{iHt} | 0 \rangle \right) \\
 &+ O((qt)^2).
 \end{aligned} \tag{22}$$

Now, the quantity $\frac{\text{Tr}[e^{iHt_1}]}{2^N}$ is of order 1 at small t_1 . Hence, the parenthesis can be positive or negative, implying noise can *amplify* the signal, counter-intuitively. Since $\text{Tr}[e^{iHt}]$ and $\langle 0 | e^{iHt} | 0 \rangle$ are concave functions of t at small t , we can expect indeed this parenthesis to be positive.

To compare this theoretical noise model to hardware results, we proceed as follows. Recall that q is the error rate per unit of physical time t , not per gate. In an actual implementation of the algorithm, the error rate q depends on the chosen gate angle, because the number of gates per unit of physical time depends on the gate angle. In the experiments, since the gate angle is set to be inversely proportional to physical time t , the error

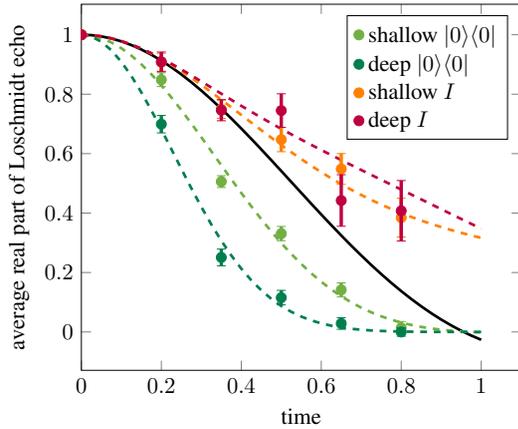


FIG. 5. Real part of Loschmidt amplitude as a function of time Jt , showing hardware results (bullets) and curves obtained from theoretical noise model (dashed lines), for shallow and deep circuits, and for measurement operator $\mathcal{O} = |0\rangle\langle 0|, I$. The single parameter entering the theoretical model is fitted with the hardware data obtained for shallow circuits and $\mathcal{O} = |0\rangle\langle 0|$ (lightgreen). The three other dashed curves are deduced from it. The black line indicates the noiseless value.

rate q is proportional to the final physical time t for each experiment. Note in the formulas above, q stays constant in time, because the gate angle is constant throughout a single time evolution. To determine the proportionality factor β between error rate q and final physical time t , we fit the formula for $\langle X \otimes \mathcal{O} \rangle_{\text{noisy}}$ in (19) for $\mathcal{O} = |0\rangle\langle 0|$, only on shallow circuits. Since the error rate must be proportional to the number of gates per unit of physical time, the error rate for deep circuits must be three times larger.

The fit yields the value $\beta = 2.46J^2$. This means that, for $Jt = 0.5$ for example, the error rate q in the shallow circuits is $\beta t = 1.23J$, and so we have on average $qt = 0.615$ errors during the evolution for time t . Since there are on average 275 TQ gates per shallow circuit for $Jt = 0.5$, this gives an average number of errors per TQ gate equal to 0.0022. The component benchmark values give an error rate per TQ gate around 10^{-3} [55]. However, this value 0.0022 obtained from the hardware also takes into account other sources of error like memory error, it is thus expected that it is larger than the component benchmark value. We plot in Fig. 5 the comparison of the theoretical noise model obtained with this fitted parameter, for the three other settings, namely shallow circuits with $\mathcal{O} = I$, and deep circuits for both $\mathcal{O} = |0\rangle\langle 0|$ and $\mathcal{O} = I$. We observe a good agreement for all the setups. In particular, the model correctly captures the upward bias due to noise of hardware results when $\mathcal{O} = I$.

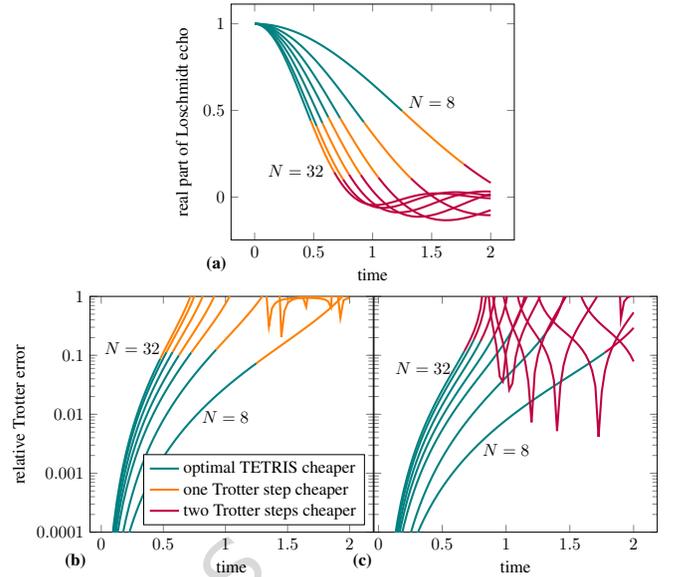


FIG. 6. **a**: real part of Loschmidt amplitude as a function of time Jt , for different system sizes $N = 8, 12, \dots, 32$ (from right to left), with green indicating that TETRIS at optimal angle is cheaper than Trotter, orange that one Trotter step is cheaper, and purple that two Trotter steps are cheaper. **b,c**: relative Trotter error in the real part of the Loschmidt amplitude, as a function of time, using one single Trotter step (**b**) and two Trotter steps (**c**), for different system sizes $N = 8, 12, \dots, 32$ (from bottom to top).

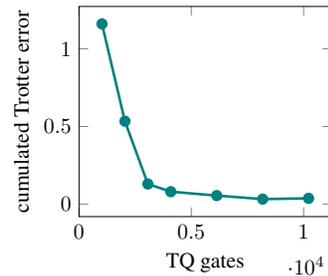


FIG. 7. Sum of the absolute value of the differences between Trotterized evolution and exact evolution of the real part of the Loschmidt echo, at time points $t = 0.1n$ for $n = 1, \dots, 20$, in size $L = 12$, as a function of the total number of TQ gates, corresponding to different number of Trotter steps, including 1000 random realizations of the SYK Hamiltonian.

Cost comparison with Trotter

We now perform a cost comparison against Trotterization. We saw that as a function of time, the complexity of our randomized algorithm scales as $O(t^2)$, whereas a Trotter decomposition at fixed Trotter step s scales as $O(t)$. In panel (a) of Fig. 6, we show, for sparsity parameter $k = 2.3$, the different regimes where TETRIS has a lower gate count than one or two Trotter steps, as a function of time. The gate count is always computed for the optimal gate angle τ . In that case, the sampling

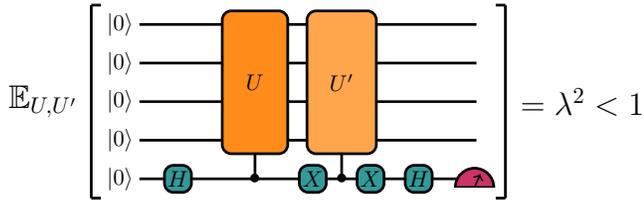


FIG. 8. Schematic circuit representation of the mirror-on-average circuit benchmark. If the unitaries U and U' are both drawn randomly and independently from the TETRIS algorithm, then in average the expectation value of the ancilla is equal to λ^2 , with λ given in (8).

overhead of TETRIS over Trotter is approximately equal to $e \approx 2.72$. We see that, at early times, TETRIS is always cheaper. Moreover, when system size increases, it remains consistently cheaper in the regime where the Loschmidt amplitude is $\gtrsim 0.5$. At later times, doing one single Trotter step always ends up getting cheaper. However, increasing time at a fixed number of Trotter steps inevitably increases the bias due to the Trotter error, which we have ignored so far.

In panels (b,c) of Fig. 6, we show the relative Trotter error as a function of time, when doing one or two Trotter steps. We also show in Fig 7 the amount of Trotter error as a function of the number of Trotter steps. We see that for either one or Trotter steps, Trotterization becomes cheaper only when the Trotter decomposition incurs at least $\approx 10\%$ error. In particular, for time < 1 , the regime where one Trotter step is cheaper than TETRIS at optimal angle would have incurred a Trotter error of the same order as the error bars on the extrapolated values in Fig. 3. Trotter error would thus have been statistically visible in Fig. 3, even neglecting hardware noise.

We note that beyond this analysis, the optimal choice between Trotter and TETRIS would depend on the tolerance of the end user to imprecision to Trotter error. If the tolerance is high, then Trotter has a better scaling for intermediate times. However, if accuracy matters, TETRIS always becomes more efficient due to bounded average gate count with precision.

Finally, let us mention that there exists a wealth of other Hamiltonian simulation techniques with better asymptotic scalings than a Trotter decomposition. This includes for example higher-order Trotterizations [57, 58], the linear combination of unitaries method [59, 60] or more recent methods [61]. While these methods are asymptotically more efficient (in time or in precision), they often involve significant overhead, which make them unlikely to rival with just a few Trotter steps for the current and near-term hardware resources. There also exists other randomized algorithms that enjoy some of the benefits of TETRIS and whose comparison would be interesting [62–66].

Noise estimation through a mirror-on-average circuit benchmark

Estimating the amount of noise in a circuit run on hardware is an important element of quantum computing calculations. However, comparing the measured expectation value to exact values is often limited to small system sizes or shallow circuits. Scalable methods to estimate the amount of noise in a circuit are to run a *mirror* circuit on top, where all the gates are exactly inverted. Since the circuit is logically equivalent to identity, hardware imperfections can be estimated in a scalable way by estimating the survival probability of the initial state [35–37]. These mirror circuit benchmarks, however, are known to overestimate the noise in actual local observables [67–69].

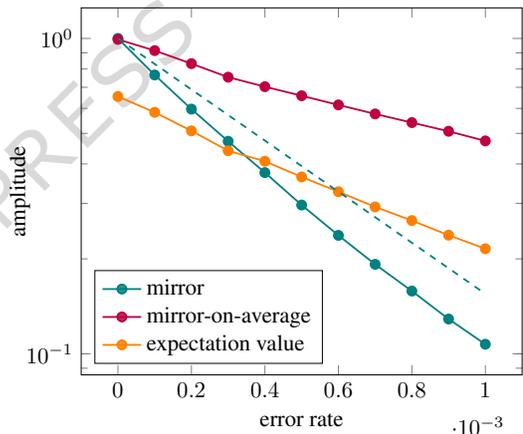


FIG. 9. Numerically simulated mirror-on-average benchmark $\lambda^{-2}\mathbb{E}_{U,U'}[\rho_{\text{ave}}(X \otimes I)]$ (24) (purple), standard mirror benchmark $\mathbb{E}_U[\langle 0|U^\dagger U|0\rangle]$ (teal), both of which return 1 in the noiseless case, and expectation values of the operator $\mathcal{O} = \frac{1}{L}\sum_j Z_j$ (orange), as a function of TQ depolarizing error rate p_{dep} , in size $L = 12$, for $Jt = 0.8$, taking disorder average of the sparse SYK Hamiltonian H . The teal dashed curve indicates the process gate fidelity, $1 - 15p_{\text{dep}}/16$, to the average number of TQ gates in the circuits, here 1988.

The TETRIS algorithm allows instead for a scalable benchmark protocol where a mirror circuit equivalent to identity is implemented only *on average*. We initialize a density matrix ρ on the ancilla and the system qubits in $|+\rangle \otimes |0\rangle$. We generate two independent random unitaries U, U' with the TETRIS algorithm, that satisfy $\mathbb{E}_U[U] = \lambda e^{iHt}$. We apply U on the system qubits conditioned on the ancilla to be in $|0\rangle$, and U' on the system qubits conditioned on the ancilla to be in $|1\rangle$. The density matrix obtained from these steps is

$$\rho = \frac{1}{2}|0\rangle\langle 0| \otimes U|0\rangle\langle 0|U^\dagger + \frac{1}{2}|0\rangle\langle 1| \otimes U|0\rangle\langle 0|U^\dagger + \frac{1}{2}|1\rangle\langle 0| \otimes U'|0\rangle\langle 0|U'^\dagger + \frac{1}{2}|1\rangle\langle 1| \otimes U'|0\rangle\langle 0|U'^\dagger. \quad (23)$$

Using (7), after averaging ρ over both U and U' we obtain

$$\begin{aligned} \rho_{\text{ave}} &:= \mathbb{E}_{U,U'}[\rho] = \\ &\frac{1}{2}|0\rangle\langle 0| \otimes \mathbb{E}_U[U|0\rangle\langle 0|U^\dagger] + \frac{\lambda^2}{2}|1\rangle\langle 0| \otimes e^{iHt}|0\rangle\langle 0|e^{-iHt} \\ &+ \frac{\lambda^2}{2}|0\rangle\langle 1| \otimes e^{iHt}|0\rangle\langle 0|e^{-iHt} + \frac{1}{2}|1\rangle\langle 1| \otimes \mathbb{E}_{U'}[U'|0\rangle\langle 0|U'^\dagger]. \end{aligned} \quad (24)$$

Similar as before, denoting $X \otimes \mathcal{O}$ the observable applying X on the ancilla and \mathcal{O} on the system qubits, we have

$$\text{Tr}[\rho_{\text{ave}}(X \otimes \mathcal{O})] = \lambda^2 \langle 0|e^{-iHt}\mathcal{O}e^{iHt}|0\rangle. \quad (25)$$

It follows that the ideal expectation value of X on the ancilla $\text{Tr}[\rho_{\text{ave}}(X \otimes I)] = \lambda^2$ is known exactly. The expectation value of X on the ancilla cannot be computed in a scalable way for individual circuits because U and U' are generally different. However, on averaging over U, U' , a *mirror-on-average* circuit is implemented, and the expectation value of X on the ancilla is exactly λ^2 . Moreover, from the same circuits, the expectation value of \mathcal{O} within $e^{iHt}|0\rangle$ can be computed. We thus expect that the noise observed on the hardware on the expectation value of X on the ancilla gives a good estimate of the noise observed on a local observable. This is in contrast with usual mirror-circuits, which measure a global observable. A schematic representation of these circuits is shown in Fig 8. Overall, the operational intuition behind the relevance of this benchmark process is that it involves rather generic circuits encountered in Hamiltonian simulation, whose individual expectation values are not known, but whose average over the random process is known for arbitrary system sizes and circuit depth. It yields thus a benchmark process that is both scalable and application-oriented, in the sense that each random circuit will be similar to some Hamiltonian simulation circuit.

In Fig. 9, we test these ideas with noisy simulations of these mirror-on-average circuits. For $L = 12$ qubits, and time $Jt = 0.8$, we plot the expectation value of $\mathcal{O} = \frac{1}{L} \sum_j Z_j$ computed this way, as well as the expectation value of X on the ancilla divided by λ^2 . We also compare with the expectation value obtained when measuring an exact mirror circuit where the inverse U^\dagger is applied on the ancilla conditioned to be 0. All the Hamiltonians are systematically averaged over different SYK realizations. We first observe that the standard mirror circuits display a decay with noise that is comparable to the circuit fidelity, estimated as gate fidelity to the number of gates (the agreement is not exact because the number of gates fluctuates among different random circuits). This is in line with previous implementations of mirror circuit benchmarks [35, 37]. Secondly, we observe that the expectation value of a local observable like $\frac{1}{L} \sum_j Z_j$ is significantly less damped by noise than this mirror circuit. Instead, its decay with noise is comparable (but

not identical) to that of the mirror-on-average circuits. This agreement between the two is not systematic, and for other parameter regimes the slope of the mirror-on-average circuits can present more deviation from than that of observables, but is always smaller than that of the mirror circuits. This shows that these mirror-on-average circuits give a better estimate of noise on actual observables than a standard mirror circuit benchmark, which tends to systematically overestimate the noise on local observables [35]. Moreover, this mirror-on-average benchmark is scalable in the number of qubits and in the circuit depth.

DISCUSSION

We have demonstrated an experimental quantum simulation of the Loschmidt amplitude for the sparse SYK model, based on the TETRIS algorithm. The TETRIS algorithm is particularly suited to simulate the SYK model because its disorder average can be readily combined with the random sampling of TETRIS circuits to reduce the sampling overhead. While the sparsification of the SYK model and the use of TETRIS significantly reduce the circuit complexity, that alone is not enough to obtain reliable estimates of physical observables by running those circuits on currently available quantum processors. To overcome this difficulty, we proposed two noise mitigation techniques: echo verification and LGAE. We extensively assessed the performance of the algorithm with the noise mitigations and demonstrated substantial improvements in the accuracy of the experimentally obtained Loschmidt amplitude.

In order to address more physically motivated questions in future experiments, we discuss the computation of an OTOC $\text{Tr}[Ae^{iHt}Be^{-iHt}A^\dagger e^{iHt}B^\dagger e^{-iHt}]$, where A and B are typically set to be Majorana operators. We note that the computation of the Lyapunov exponent requires the calculation of OTOCs at finite temperature. Preparation of the thermal states at scale requires separate consideration and is left for future work. Here, we replace the trace by the average of the expectation values with respect to randomly sampled pure states. With the TETRIS algorithm and the ternary tree encoding [41], which maps a fermion operator to a spin operator of weight $[\log_3 2L]$, each time-evolution operator $e^{\pm iHt}$ uses approximately $t^2 \mu^2 \log_3(2L) \approx 2k(Jt)^2 L^2 \log_3(2L)$ TQ gates. This TQ gate count assumes using standard Pauli gadgets to implement rotations of Pauli strings with several Pauli matrices. The number of non-Clifford gates is thus $\sim t^2 \mu^2 \approx 2k(Jt)^2 L^2$. The conversion of this number to the number of T -gates depends on the precision of the rotation angles. See e.g. [70]. To extract the quantum Lyapunov exponent from the decay of the OTOC, we set the simulation time to $Jt \sim \ln N$. Therefore, the total

TQ gate count is roughly,

$$8 \times k(\ln(2L)L)^2 \log_3(2L) \approx \begin{cases} 4 \times 10^6 & (L = 50) \\ 2 \times 10^7 & (L = 100) \end{cases} \quad (26)$$

for $k = 2.3$. Without any parallelization of gates and assuming the execution time per circuit depth being 30 milliseconds (ms), the runtime of each circuit is roughly $4 \times 10^6 \times 30 \text{ ms} = 30$ hours for $L = 50$. Since $L/\log_3(L) \approx 14$ gates can be parallelized by adding the same number of ancillary qubits, the execution time per circuit is reduced to two hours.

These estimates point towards the need for further algorithmic improvements, for example, by using techniques like sum-of-squares spectral amplification [71] or defining new simplified Hamiltonians that retain the chaotic properties of the original dense SYK model (see e.g. [72–74] for recent explorations of SYK-like models).

METHODS

Noise mitigation: Large Gate Angle Extrapolation (LGAE)

In TETRIS, the gate angle τ is a free parameter that does not modify the precision obtained on the result: the real part of the Loschmidt amplitude can be obtained with Eq. (13) for all τ . Yet, the gate angle τ has a significant impact on the unitaries U sampled in Eq. (7) and on the attenuation factor λ in Eq. (8). We leverage this freedom to introduce another error mitigation scheme. The angle of the rotations in a TETRIS circuit U is τ , and there are on average $t\mu/\sin\tau$ such rotations in each circuit U . Being able to tune a parameter to continuously change the number of gates in the circuit, but without changing the expectation values, is the ideal setup to perform Zero Noise Extrapolation (ZNE) [75, 76]. Namely, on an actual noisy hardware, by changing τ one can measure the effect of noise on the result, and compensate for it. We refer to this variant of ZNE as Large Gate Angle Extrapolation (LGAE), in the sense that extrapolating noise to zero in the ZNE is performed here by increasing the gate angle.

LGAE proceeds as follows. We select two values of τ , a fixed $\tau_0 > 0$ and a rescaled $\tau_\alpha = \alpha\tau_0$, where $0 < \alpha < 1$. Then, we run TETRIS with both angles and let y_0 and y_α be the expectation values obtained from the corresponding circuits. Let us assume for simplicity that these gate angles are sufficiently small so that $\sin\tau \approx \tau$. Also, as usual with ZNE, we assume that, at low noise level, the effect of noise is linear in the number of gates,

$$y_\alpha = a + bN_g(\alpha), \quad (27)$$

with the parameters a, b and the average number of gates $N_g(\alpha)$ in the circuits generated with gate angle τ_α . This

number of gates is proportional to $1/\sin\tau_\alpha \approx 1/\tau_\alpha$. Under these assumptions, we obtain a mitigated expectation value as

$$y_{\text{LGAE}} = \frac{y_0 - \alpha y_\alpha}{1 - \alpha}. \quad (28)$$

The standard deviation σ_{LGAE} obtained on y_{LGAE} is

$$\sigma_{\text{LGAE}} = \frac{1}{1 - \alpha} \sqrt{\sigma_0^2 + \alpha^2 \sigma_\alpha^2}, \quad (29)$$

where σ_0 and σ_α are the standard deviations of y_0, y_α , respectively. Considering that there is a certain total number of shots to distribute among σ_0 and σ_α , we write $\sigma_0^2 = s_0^2/x$ and $\sigma_\alpha^2 = s_\alpha^2/(1-x)$, with $0 < x < 1$ and s_0, s_α numbers. The minimal variance is obtained for

$$x = \frac{s_0}{s_0 + \alpha s_\alpha}, \quad (30)$$

and then it takes the value

$$\sigma_{\text{LGAE}} = \frac{s_0 + \alpha s_\alpha}{1 - \alpha}. \quad (31)$$

Similarly to the conventional ZNE, instead of the linear fit with the number of gates, we may employ an *exponential* fit. The assumption is now that the effect of noise is of the form

$$y_\alpha = ae^{bN_g(\alpha)}. \quad (32)$$

In that case the exponential mitigation gives

$$y_{\text{exp-LGAE}} = \exp \frac{\log y_0 - \alpha \log y_\alpha}{1 - \alpha}. \quad (33)$$

There is no closed-form expression for the standard deviation of this estimate in terms of σ_0, σ_α in general, but assuming these are small, we can write using Gaussian propagation of uncertainty

$$\sigma_{\text{exp-LGAE}} = \frac{y_{\text{exp-LGAE}}}{1 - \alpha} \sqrt{\frac{\alpha^2 \sigma_\alpha^2}{y_\alpha^2} + \frac{\sigma_0^2}{y_0^2}}. \quad (34)$$

In panel (b) of Fig. 2, we show the measured real part of the Loschmidt amplitude for different values of gate angles, when performing noisy numerical simulations of the SYK model for $N = 24$ fermions, and for the three operators $\mathcal{O} = I, |0\rangle\langle 0|, |0\rangle\langle 0|_{\text{mit}}$ discussed in the previous section. When running noiseless simulations, all the points would match the exact value, independently of the gate angle. However, in noisy simulations, we see that the signal measured for $\mathcal{O} = |0\rangle\langle 0|, |0\rangle\langle 0|_{\text{mit}}$ depends strongly on the gate angle, and decays when the gate angle decreases. The signal for $\mathcal{O} = |0\rangle\langle 0|_{\text{mit}}$ is approximately linear, while that for $\mathcal{O} = |0\rangle\langle 0|$ presents a curvature. The signal for $\mathcal{O} = I$ is much less sensitive to gate angle, and is even slightly larger than the exact value due to the noise effect that we will discuss later (see e.g. Fig. 5). Because of the peculiar non-monotonous effect of noise here, the noise mitigation works less well. The error bars show that $\mathcal{O} = I$ displays more shot noise, requiring around twice as many shots to get the same precision as for $\mathcal{O} = |0\rangle\langle 0|, |0\rangle\langle 0|_{\text{mit}}$.

ACKNOWLEDGEMENTS

The experimental data reported in this work were produced by the Quantinuum H1-1 quantum computer, Powered by Honeywell, on February 6-18, 2025. E.G. acknowledges support by the Bavarian Ministry of Economic Affairs, Regional Development and Energy (StMWi) under project Bench-QC (DIK0425/01). Y.K. thanks Juan Pedersen for useful discussions. We thank Matthew DeCross and Christopher Self for reading the paper and providing feedback.

DATA AVAILABILITY STATEMENT

The hardware data generated during this study is available upon request.

COMPETING INTERESTS

H.D. is a shareholder of Quantinuum.

AUTHOR CONTRIBUTIONS

Y.K. and E.R. initiated the idea to use the stochastic algorithm for simulating the SYK model, did the background research and scaling analysis. E.G. implemented the numerical analysis, the hardware experiments and made the figures. E.G. and H.D. developed the noise mitigation technique. All authors contributed to discussing the results and improving the hardware protocol. All authors contributed to writing the manuscript.

REFERENCES

-
- * etienne.granet@quantinuum.com
- [1] Sachdev, S. & Ye, J. Gapless spin-fluid ground state in a random quantum heisenberg magnet. *Phys. Rev. Lett.* **70**, 3339–3342 (1993).
 - [2] Kitaev, A. A simple model of quantum holography (2015). URL <https://online.kitp.ucsb.edu/online/entangled15/>.
 - [3] Maldacena, J. & Stanford, D. Remarks on the Sachdev-Ye-Kitaev model. *Phys. Rev. D* **94**, 106002 (2016). 1604.07818.
 - [4] Luo, Z. *et al.* Quantum simulation of the non-fermi-liquid state of Sachdev-Ye-Kitaev model. *npj Quantum Inf.* **5**, 53 (2019). 1712.06458.
 - [5] Chowdhury, D., Georges, A., Parcollet, O. & Sachdev, S. Sachdev-Ye-Kitaev models and beyond: Window into non-Fermi liquids. *Rev. Mod. Phys.* **94**, 035004 (2022). 2109.05037.
 - [6] Sachdev, S. Statistical mechanics of strange metals and black holes. *ICTS News* **8** (2022). 2205.02285.
 - [7] Danshita, I., Hanada, M. & Tezuka, M. Creating and probing the Sachdev-Ye-Kitaev model with ultracold gases: Towards experimental studies of quantum gravity. *PTEP* **2017**, 083I01 (2017). 1606.02454.
 - [8] Schuster, T. *et al.* Many-body quantum teleportation via operator spreading in the traversable wormhole protocol. *Phys. Rev. X* **12**, 031013 (2022).
 - [9] Brown, A. R. *et al.* Quantum gravity in the lab. i. teleportation by size and traversable wormholes. *PRX Quantum* **4**, 010320 (2023).
 - [10] Nezami, S. *et al.* Quantum gravity in the lab. ii. teleportation by size and traversable wormholes. *PRX Quantum* **4**, 010321 (2023).
 - [11] Shapoval, I., Su, V. P., de Jong, W., Urbanek, M. & Swingle, B. Towards Quantum Gravity in the Lab on Quantum Processors. *Quantum* **7**, 1138 (2023). 2205.14081.
 - [12] Sárosi, G. AdS₂ holography and the SYK model. *PoS Modave2017*, 001 (2018). 1711.08482.
 - [13] Rosenhaus, V. An introduction to the SYK model. *J. Phys. A* **52**, 323001 (2019). 1807.03334.
 - [14] Maldacena, J., Shenker, S. H. & Stanford, D. A bound on chaos. *JHEP* **08**, 106 (2016). 1503.01409.
 - [15] Jafferis, D. *et al.* Traversable wormhole dynamics on a quantum processor. *Nature* **612**, 51–55 (2022).
 - [16] Asaduzzaman, M., Jha, R. G. & Sambasivam, B. Sachdev-ye-kitaev model on a noisy quantum computer. *Physical Review D* **109**, 105002 (2024).
 - [17] Xu, S., Susskind, L., Su, Y. & Swingle, B. A Sparse Model of Quantum Holography (2020). 2008.02303.
 - [18] García-García, A. M., Jia, Y., Rosa, D. & Verbaarschot, J. J. M. Sparse sachdev-ye-kitaev model, quantum chaos, and gravity duals. *Phys. Rev. D* **103**, 106002 (2021).
 - [19] Cáceres, E., Misobuchi, A. & Pimentel, R. Sparse SYK and traversable wormholes. *JHEP* **11**, 015 (2021). 2108.08808.
 - [20] Cáceres, E., Misobuchi, A. & Raz, A. Spectral form factor in sparse SYK models. *JHEP* **08**, 236 (2022). 2204.07194.
 - [21] Cáceres, E., Guglielmo, T., Kent, B. & Misobuchi, A. Out-of-time-order correlators and Lyapunov exponents in sparse SYK. *JHEP* **11**, 088 (2023). 2306.07345.
 - [22] García-García, A. M., Liu, C. & Verbaarschot, J. J. M. Sparsity-Independent Lyapunov Exponent in the Sachdev-Ye-Kitaev Model. *Phys. Rev. Lett.* **133**, 091602 (2024). 2311.00639.
 - [23] Orman, P., Gharibyan, H. & Preskill, J. Quantum chaos in the sparse SYK model. *JHEP* **02**, 173 (2025). 2403.13884.
 - [24] Granet, E. & Dreyer, H. Hamiltonian dynamics on digital quantum computers without discretization error. *npj Quantum Information* **10**, 82 (2024).
 - [25] Peres, A. Stability of quantum motion in chaotic and regular systems. *Phys. Rev. A* **30**, 1610–1615 (1984).
 - [26] Goussev, A., Jalabert, R. A., Pastawski, H. M. & Wisniacki, D. Loschmidt echo (2012). 1206.6348.
 - [27] Gorin, T., Prosen, T., Seligman, T. H. & Žnidarič, M. Dynamics of Loschmidt echoes and fidelity decay. *Phys. Rept.* **435**, 33–156 (2006).

- [28] Yan, B., Cincio, L. & Zurek, W. H. Information Scrambling and Loschmidt Echo. *Phys. Rev. Lett.* **124**, 160603 (2020). 1903.02651.
- [29] Larkin, A. I. & Ovchinnikov, Y. N. Quasiclassical Method in the Theory of Superconductivity. *Soviet Journal of Experimental and Theoretical Physics* **28**, 1200 (1969).
- [30] Shenker, S. H. & Stanford, D. Black holes and the butterfly effect. *Journal of High Energy Physics* **2014**, 067–067 (2014).
- [31] Roberts, D. A., Stanford, D. & Susskind, L. Localized shocks. *Journal of High Energy Physics* **2015**, 051 (2015).
- [32] Hosur, P., Qi, X.-L., Roberts, D. A. & Yoshida, B. Chaos in quantum channels. *Journal of High Energy Physics* **2016**, 004 (2016).
- [33] Swingle, B. Unscrambling the physics of out-of-time-order correlators. *Nature Phys.* **14**, 988–990 (2018).
- [34] Xu, S. & Swingle, B. Scrambling Dynamics and Out-of-Time-Ordered Correlators in Quantum Many-Body Systems. *PRX Quantum* **5**, 010201 (2024).
- [35] Proctor, T., Rudinger, K., Young, K., Nielsen, E. & Blume-Kohout, R. Measuring the capabilities of quantum computers. *Nature Phys.* **18**, 75–79 (2022). 2008.11294.
- [36] Mayer, K. *et al.* Theory of mirror benchmarking and demonstration on a quantum computer (2021). 2108.10431.
- [37] DeCross, M. *et al.* The computational power of random quantum circuits in arbitrary geometries. *arXiv preprint arXiv:2406.02501* (2024).
- [38] Granet, E. & Dreyer, H. Appqsim: Application-oriented benchmarks for hamiltonian simulation on a quantum computer. *arXiv preprint arXiv:2503.04298* (2025).
- [39] Derby, C., Klassen, J., Bausch, J. & Cubitt, T. Compact fermion to qubit mappings. *Physical Review B* **104**, 035118 (2021).
- [40] Nigmatullin, R. *et al.* Experimental demonstration of breakeven for a compact fermionic encoding. *Nature Physics* **1–7** (2025).
- [41] Jiang, Z., Kalev, A., Mruczkiewicz, W. & Neven, H. Optimal fermion-to-qubit mapping via ternary trees with applications to reduced quantum states learning. *Quantum* **4**, 276 (2020).
- [42] Chen, Y., Helsen, J. & Ozols, M. Trotter error and gate complexity of the SYK and sparse SYK models (2025). 2502.18420.
- [43] Campbell, E. Random Compiler for Fast Hamiltonian Simulation. *Phys. Rev. Lett.* **123**, 070503 (2019).
- [44] Childs, A. M., Ostrander, A. & Su, Y. Faster quantum simulation by randomization. *Quantum* **3**, 182 (2019).
- [45] Chen, C.-F., Huang, H.-Y., Kueng, R. & Tropp, J. A. Concentration for random product formulas. *PRX Quantum* **2**, 040305 (2021).
- [46] Wan, K., Berta, M. & Campbell, E. T. Randomized Quantum Algorithm for Statistical Phase Estimation. *Phys. Rev. Lett.* **129**, 030503 (2022). 2110.12071.
- [47] Nakaji, K., Bagherimehrab, M. & Aspuru-Guzik, A. High-Order Randomized Compiler for Hamiltonian Simulation. *PRX Quantum* **5**, 020330 (2024). 2302.14811.
- [48] Pocrnic, M., Hagan, M., Carrasquilla, J., Segal, D. & Wiebe, N. Composite Qdrift-product formulas for quantum and classical simulations in real and imaginary time. *Phys. Rev. Res.* **6**, 013224 (2024). 2306.16572.
- [49] Watson, J. D. Randomly Compiled Quantum Simulation with Exponentially Reduced Circuit Depths (2024). 2411.04240.
- [50] Kiumi, C. & Koczor, B. TE-PAI: exact time evolution by sampling random circuits. *Quantum Sci. Technol.* **10**, 045071 (2025). 2410.16850.
- [51] O’Brien, T. E. *et al.* Error mitigation via verified phase estimation. *PRX Quantum* **2**, 020317 (2021).
- [52] Huo, M. & Li, Y. Dual-state purification for practical quantum error mitigation. *Phys. Rev. A* **105**, 022427 (2022).
- [53] Cai, Z. Resource-efficient Purification-based Quantum Error Mitigation (2021). 2107.07279.
- [54] Polla, S., Anselmetti, G.-L. R. & O’Brien, T. E. Optimizing the information extracted by a single qubit measurement. *Phys. Rev. A* **108**, 012403 (2023).
- [55] Quantinuum. H1-1 product data sheet (2025). URL https://docs.quantinuum.com/systems/data_sheets/Quantinuum%20H1%20Product%20Data%20Sheet.pdf.
- [56] Haghsheenas, R. *et al.* Digital quantum magnetism at the frontier of classical simulations. *arXiv preprint arXiv:2503.20870* (2025).
- [57] Suzuki, M. General theory of fractal path integrals with applications to many-body theories and statistical physics. *Journal of mathematical physics* **32**, 400–407 (1991).
- [58] Childs, A. M., Su, Y., Tran, M. C., Wiebe, N. & Zhu, S. Theory of trotter error with commutator scaling. *Physical Review X* **11**, 011020 (2021).
- [59] Berry, D. W., Childs, A. M., Cleve, R., Kothari, R. & Somma, R. D. Simulating hamiltonian dynamics with a truncated taylor series. *Physical review letters* **114**, 090502 (2015).
- [60] Childs, A. M. & Wiebe, N. Hamiltonian simulation using linear combinations of unitary operations. *arXiv preprint arXiv:1202.5822* (2012).
- [61] Watson, J. D. & Watkins, J. Exponentially reduced circuit depths using trotter error mitigation. *PRX Quantum* **6**, 030325 (2025).
- [62] Campbell, E. Random compiler for fast hamiltonian simulation. *Phys. Rev. Lett.* **123**, 070503 (2019).
- [63] Ouyang, Y., White, D. R. & Campbell, E. T. Compilation by stochastic hamiltonian sparsification. *Quantum* **4**, 235 (2020).
- [64] Chen, C.-F., Huang, H.-Y., Kueng, R. & Tropp, J. A. Concentration for random product formulas. *PRX Quantum* **2**, 040305 (2021).
- [65] Nakaji, K., Bagherimehrab, M. & Aspuru-Guzik, A. High-order randomized compiler for hamiltonian simulation. *PRX Quantum* **5**, 020330 (2024).
- [66] Zeng, P., Sun, J., Jiang, L. & Zhao, Q. Simple and high-precision hamiltonian simulation by compensating trotter error with linear combination of unitary operations. *PRX Quantum* **6**, 010359 (2025).
- [67] Schiffer, B. F., Rubio, A. F., Trivedi, R. & Cirac, J. I. The quantum adiabatic algorithm suppresses the proliferation of errors. *arXiv preprint arXiv:2404.15397* (2024).
- [68] Granet, E. & Dreyer, H. Dilution of error in digital hamiltonian simulation. *PRX Quantum* **6**, 010333 (2025).
- [69] Chertkov, E., Chen, Y.-H., Lubasch, M., Hayes, D. & Foss-Feig, M. Robustness of near-thermal dynamics on digital quantum computers. *arXiv preprint arXiv:2410.10794* (2024).

- [70] Kliuchnikov, V. & Lauter, K. & Minko, R. & Paetznick, A. & Petit, C. Shorter quantum circuits via single-qubit gate approximation. *Quantum* **7**, 1208 (2023).
- [71] King, R., Low, G. H., Babbush, R., Somma, R. D. & Rubin, N. C. Quantum simulation with sum-of-squares spectral amplification (2025). 2505.01528.
- [72] Swingle, B. & Winer, M. Bosonic model of quantum holography. *Phys. Rev. B* **109**, 094206 (2024). 2311.01516.
- [73] Hanada, M., Jevicki, A., Liu, X., Rinaldi, E. & Tezuka, M. A model of randomly-coupled Pauli spins. *JHEP* **05**, 280 (2024). 2309.15349.
- [74] Hanada, M., van Leuven, S., Oktay, O. & Tezuka, M. Two-local modifications of SYK model with quantum chaos (2025). 2505.09900.
- [75] Temme, K., Bravyi, S. & Gambetta, J. M. Error Mitigation for Short-Depth Quantum Circuits. *Phys. Rev. Lett.* **119**, 180509 (2017). 1612.02058.
- [76] Giurgica-Tiron, T., Hindy, Y., LaRose, R., Mari, A. & Zeng, W. J. Digital zero noise extrapolation for quantum error mitigation. In *2020 IEEE International Conference on Quantum Computing and Engineering* (2020). 2005.10921.

ARTICLE IN PRESS