

Integer factorization via tensor-network Schnorr's sieving

Marco Tesoro ^{1,2,3} Ilaria Siloi ^{1,2,3} Daniel Jaschke ^{1,2,3,4} Giuseppe Magnifico ^{1,2,3,5} and Simone Montangero ^{1,2,3}

¹*Dipartimento di Fisica e Astronomia “G. Galilei”, Università degli Studi di Padova, via Marzolo 8, I-35131 Padova, Italy*

²*Padua Quantum Technologies Research Center, Università degli Studi di Padova, Padova, Italy*

³*INFN, Sezioni di Padova & Bari, Padova, Italy*

⁴*Institute for Complex Quantum Systems, Ulm University, Albert-Einstein-Allee 11, 89069 Ulm, Germany*

⁵*Dipartimento di Fisica, Università di Bari, I-70126 Bari, Italy*



(Received 17 December 2024; accepted 22 January 2026; published 11 March 2026)

Classical public-key cryptography standards rely on the Rivest-Shamir-Adleman (RSA) encryption protocol. The security of this protocol is based on the exponential computational complexity of the most efficient classical algorithms for factoring large semiprime numbers into their two prime components. Here, we address RSA factorization building on Schnorr's mathematical framework where factorization translates into a combinatorial optimization problem. We solve the optimization task via tensor network methods, a quantum-inspired classical numerical technique. This tensor network Schnorr's sieving algorithm displays numerical evidence of polynomial scaling of resources with the bit-length of the semiprime. We factorize RSA numbers up to 100 bits and assess how computational resources scale through numerical simulations up to 130 bits, encoding the optimization problem in quantum systems with up to 256 qubits. Only the high-order polynomial scaling of the required resources limits the factorization of larger numbers. Although these results do not currently undermine the security of the present communication infrastructure, they strongly highlight the urgency of implementing postquantum cryptography or quantum key distribution.

DOI: [10.1103/d9dl-ctt4](https://doi.org/10.1103/d9dl-ctt4)

I. INTRODUCTION

Among modern encryption schemes, the Rivest-Shamir-Adleman (RSA) protocol stands out for its resilience against cyber threats from conventional computers [1]. The security of RSA encryption is rooted in the computational complexity of factoring any large semiprime N in its two prime components p and q , i.e., $N = pq$ [2,3]. State-of-the-art factoring algorithms for RSA numbers are sieving methods [4–6], which rely on the observation that any odd number can be expressed as the difference of two squares $N = (\frac{p+q}{2})^2 - (\frac{p-q}{2})^2 = X^2 - Y^2$, where X and Y are integers. The search for prime factors then revolves around finding nontrivial solutions of the congruence relation $X^2 \equiv Y^2 \pmod{N}$. Sieving techniques differ in how potential values of X and Y are obtained. Then, it is efficient to compute the corresponding p and q and check if they factorize N [4].

The algorithmic complexity of sieving approaches has improved over the years through the application of sophisticated methods for constructing congruences of squares [7]. The most efficient known classical algorithm is the general number field sieve (GNFS), which exhibits a subexponential algorithmic complexity [8,9] and has led to the successful factorization of the largest recorded RSA number (829 bits) in 2020 [10]. Shor's quantum factorization algorithm constructs

solutions to congruence relations [11] in polynomial time, marking a transformative advancement in the field. However, attacking state-of-the-art RSA remains out of reach for current near-term quantum devices [12]. In this work, we focus on Schnorr's lattice sieving [13,18], a classical algorithm that recasts RSA factorization as the search for optimal solutions to a set of NP-hard optimization problems on a lattice, known as closest vector problems (CVPs) [19,20]. We improve Schnorr's lattice sieving by introducing an approach based on tensor network (TN) methods [21,22]. Following Yan *et al.* [15], CVP solutions are encoded in the low-energy configurations of a spin-glass Hamiltonian. We employ an optimal sampling strategy based on TN to efficiently access these configurations [17]. Finally, we construct the two squares X^2 and Y^2 satisfying the congruence relation from the combination of these configurations.

The TN Schnorr's sieving (TNSS) allows the factorization of RSA semiprimes up to over 100 bits. Moreover, we unveil how to effectively choose the hyperparameters of the algorithm for successful factorization with any key bit-length ℓ [15,23–25]. We provide a systematic numerical analysis of the resources needed by the TNSS algorithm up to 130 bits using up to 256 qubits, which scale polynomially with the key bit-length. Our findings, unless disproved by larger-scale simulations or mathematical proof, suggest that factoring RSA keys requires polynomial classical resources using the TNSS algorithm.

A. Tensor network Schnorr's sieving factoring algorithm

We propose the TNSS factorization algorithm as sketched in Fig. 1. Our scheme builds on Schnorr's idea of using

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Sketch of the TNSS algorithm. The factorization of the RSA number $N = pq$ into its prime components is mapped into an optimization problem on a lattice (lattice problem mapping). This problem is initially approached by searching for approximate solutions to a series of N -related CVPs, each defined by a lattice-target pair $(\Lambda, \mathbf{t})$ in $n+1$ dimensions [13]. A representative of the CVP set for the simplest case $n=1$ is depicted. The lattice basis $\mathcal{B}$ contains only one vector $\mathbf{b}_1 \in \mathbb{Z}^2$, which generates the infinite discrete set of lattice points $\mathbf{b} \in \Lambda$ (black points). The blue point is the target point $\mathbf{t} \in \mathbb{Z}^2$ defining the CVP $(\Lambda, \mathbf{t})$. A single approximate solution $\mathbf{b}^{cl} \in \Lambda$ (orange point) to the CVP $(\Lambda, \mathbf{t})$ is provided by Babai's classical algorithm [14]. To improve over the classical solution, we look for additional lattice points potentially encoding sr-pairs in the eigenstates of a spin glass Hamiltonian describing an n -qubit system (spin glass mapping) [15]. The 2^n configurations correspond to the set of lattice points (light gray area) around Babai's solution. The form of the Hamiltonian is defined in Eq. (1). To compute the eigenstate of the spin glass Hamiltonian, we exploit TN methods (tree tensor network sieving). By variationally minimizing the energy, we obtain a tree tensor network (TTN) state representing a superposition of low-energy eigenstates of $\mathcal{H}$ [16]. Finally, we sample a sufficient number of classical eigenstate configurations and extract their probabilities [17]. We discard all states violating the sr-pairs condition on the smoothness basis P_2 (gray lattice points) and collect the sr-pairs (green lattice points) from each CVP. These pairs are then combined to compute the two prime components p and q of the input RSA number N in the final processing step.

lattice-based sieving to filter candidate congruence relations for factoring the semiprime N (see Appendix A 1). The first phase of Schnorr's algorithm leads to the collection of a pool of candidates by searching for approximate solutions to the set of CVPs associated with N . In the second phase, further processing via linear algebra outputs the final prime components p and q [9,13,18,26]. We leverage TNs to improve the efficiency of the collection phase, being exponentially time consuming in the original formulation of the algorithm (see the discussion in Appendix A 4).

Definitions. Given a semiprime N with bit-length $\ell = \lceil \log_2 N + 1 \rceil$, we define a list of consecutive prime numbers, i.e., a *factoring basis*, for decomposing N : $P_1 = \{p_1, p_2, \dots, p_{\pi_1}\}$, where $\pi_1 = n$ is the size of P_1 . Following the procedure outlined in Appendix A 1, we construct a set of CVPs associated with N . Each CVP considers an n -rank lattice Λ described by the lattice basis $\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n\}$, where $\mathbf{b}_j \in \mathbb{Z}^{n+1}$ for $j = 1, \dots, n$. Further, we introduce a target vector $\mathbf{t} \in \mathbb{Z}^{n+1}$ which does not belong to the lattice. The lattice-target pair $(\Lambda, \mathbf{t})$ defines the CVP, where we search for the closest lattice point $\mathbf{b} \in \Lambda$ to the target $\mathbf{t}$ [2].

The above construction associates to each lattice point a pair of integers (see Appendix A 1). For a lattice point to be useful for factorization, each pair of integers must abide by the smoothness property: their ratio must not include prime factors larger than a chosen *smoothness bound* $B_2 = p_{\pi_2}$. We define the *smoothness basis*, i.e., a prime number basis P_2 of size π_2 . Each pair, referred to as a smooth-relation (sr) pair, corresponds to a lattice point sufficiently close to the target $\mathbf{t}$ [13]. Mathematically, each sr-pair represents a congruence relation between smooth integers modulo N .

Sieving, the collection phase. Babai's nearest plane algorithm [14,27] provides a single solution $\mathbf{b}^{cl} \in \Lambda$ to the CVP, representing an approximation to the closest lattice point to the target. For those CVPs where $\mathbf{b}^{cl}$ completely factorizes over the chosen smoothness basis P_2 , Babai's algorithm finds a sr-pair in polynomial time. To increase the pool of approximate solutions of a single CVP on Λ and, thus, the number of sr-pairs, one constructs an optimization problem [15] that minimizes the Euclidean distance between the target $\mathbf{t}$ and lattice points around $\mathbf{b}^{cl}$. The corresponding all-to-all n -qubit

Hamiltonian [15] reads

$$\mathcal{H} = \sum_{k=1}^{n+1} \left[t_k - b_k^{cl} - \sum_{j=1}^n \text{sign}(\mu_j - c_j) d_{j,k} Z_j \right]^2, \quad (1)$$

where $Z_j = \frac{1}{2}(\mathbb{1} - \sigma_j^z)$ and σ_j^z is the z -Pauli matrix. Each vector $\mathbf{d}_j \in \mathbb{Z}^{n+1}$ for $j = 1, \dots, n$ is obtained by reducing the lattice basis $\mathcal{B}$ of Λ via the Lenstra-Lenstra-Lovász (LLL) reduction [28], while $\{\mu_j\}_{j=1}^n$ are the Gram-Schmidt coefficients from the orthogonalization of $\mathcal{B}$. The values of $\{c_j\}_{j=1}^n$ are derived by rounding each μ_j to the nearest integer: this is a fundamental approximation in Babai's method [2]. As $\mathcal{H}$ is diagonal in the computational basis, each eigenstate corresponds to one of the 2^n possible roundings when constructing a new approximate solution around the classical one. Consequently, the eigenstates of $\mathcal{H}$ encode additional lattice points and eventually lead to sr-pairs useful for factoring N . Moreover, the energy of each eigenstate gives the square of the distance between the associated lattice point and the target $\mathbf{t}$.

Heuristic arguments associate low-energy eigenstates with lattice points that may have a larger probability of being sr-pairs [13,18]. To improve the efficiency of the collection phase, we search for these states via TN optimization: approximating the wave function as a tree tensor network (TTN) [16,22,29] and using a variational ground-state search algorithm, we find a TTN state with sufficient overlap with low-energy eigenstates (see Appendixes A 2 and A 3). The lattice points are then extracted using an efficient sampling algorithm [17]. Both algorithms are implemented in the quantum tea suite [30,31].

Processing phase. Upon gathering at least $\pi_2 + 2$ independent sr-pairs, we apply a linear algebra step to pinpoint the subset of pairs to combine, as detailed in Appendix A 1. From their combination, we obtain the two congruent squares X^2 and Y^2 for the derivation of p and q prime factors. This postprocessing step is efficiently implemented by leveraging index calculus and linear algebra techniques [32].

B. Addressing the curse of sr-pairs

Collecting enough congruence relations between smooth numbers (sr-pairs) is the only requirement for successful factorization [33]. The number of qubits n used to encode the RSA key, i.e., the lattice rank, and the smoothness bound $B_2 = p_{\pi_2}$ are tunable hyperparameters of the algorithm. Schnorr's original prescription [13,15] suggests increasing both n and π_2 sublinearly with the bit-length, $n = \pi_2 = \ell / \log_2 \ell$. However, these choices prohibit the factorization of semiprimes larger than $\ell = 70$ bits, in practice [23–25,34], due to the exponentially vanishing number of sr-pairs extracted from sublinear dimensional CVPs. For more details, we refer to the discussion in Appendix A 4. By leveraging TNs, we relax these conditions: to target more sr-pairs, we choose the smoothness bound to increase polynomially with the bit-length and the number of qubits, i.e., $\pi_2 = 2n\ell$, and explore lattice sizes beyond Schnorr's sublinear prescription (up to $n = 256$ qubits). Further, a polynomial number of bit-strings $\mathcal{O}(\ell^\gamma)$ is sampled from the TTN state, where γ is an integer hyperparameter controlling the number of low-energy configurations extracted

from the single CVP. The TNSS algorithm successfully factorizes up to 100-bit RSA keys using $n = 64$ qubits and $\gamma = 2$. Details of the 100-bit factorization are reported in Appendix A 5.

Choosing the size of the lattice according to the semiprime bit-length is crucial to ensure that enough sr-pairs are collected, once π_2 is set. To this aim, we consider bit-lengths up to $\ell = 130$ bits, and, for each semiprime, we explore lattices of different sizes. For each n , we generate $n_{\text{CVP}} = 50$ CVPs, then sample a polynomial number of bit-strings from the TTN state. Figure 2(a) shows that the *average number of sr-pair per lattice* (AsrPL) ρ_{sr} obeys

$$\rho_{sr}(\ell, n, \gamma) \sim C_1 \ell^\gamma \exp \left\{ -C_2 \left[\frac{\ell}{n^{1/\omega}} \right]^\mu \right\}, \quad (2)$$

where $C_1, C_2, \mu \in \mathbb{R}$ are fitting parameters and $\omega \approx 8.3$. Each data point in Fig. 2(a) is the average over ten semiprimes randomly drawn from an RSA key generator [36]. The value of ω and the validity of the statistical model were verified through a bootstrapped chi-squared test and a collapse quality score analysis. To discern between a polynomial and a power-law scaling, we performed additional checks to confirm the distinguishability of the two models, as detailed in Appendix A 6. We emphasize that the scaling of the number of qubits in Eq. (3) obeys a power-law function of the input bit-length ℓ , and not a high-order polynomial, making the distinction between the two models statistically resolvable within the simulated range.

The scaling of the average number of sr-pairs per lattice (AsrPL) shows again an exponential decay. This behavior is consistent with theoretical expectations from number theory, specifically, the Dickman function, which describes the distribution of smooth numbers over integers [4,25]. However, as a result of the chosen hyperparameters, Eq. (2) highlights a dependence on the effective bit-length, $\ell / n^{1/\omega}$. Finally, the scaling for the number of qubits n reads

$$n(\ell, \rho_{sr}, \gamma) \sim \left[\frac{C_2 \ell^\mu}{\log C_1 - \log \rho_{sr} + \gamma \log \ell} \right]^\mu, \quad (3)$$

where the size of the Hilbert space constraints the number of sampled bit-strings to the order $\gamma \leq n \log 2 / \log \ell$.

Figure 2(b) shows the polynomial scaling of the number of qubits with the key size when sampling a polynomial number of bit-strings as dictated by Eq. (3), assuming to find at least one sr-pair for each lattice, i.e., $\rho_{sr} = 1$. The more samples we collect, the smaller the required lattice size. The presence of at least one sr-pair per lattice allows targeting RSA keys within a polynomial number $n_{\text{CVP}} = \mathcal{O}(2n\ell)$ of CVPs. This guarantees that at least one nontrivial combination of sr-pairs forms a congruence of squares modulo N , see Appendix A 1. Although the resources needed to factorize state-of-the-art RSA keys are still beyond current capabilities, these numerical findings indicate that the resources required by the TNSS scale polynomially with the input bit-length. In Eq. (7) of Sec. ID, we provide the scaling of the number of operations for the entire factoring procedure that follows from Eq. (3).




FIG. 2. Scaling of the relevant resources in the TNSS algorithm. (a) The rescaled average number of sr-pairs per lattice (AsrPL) $\rho_{sr}^{\text{eff}} = \rho_{sr} / \ell^\gamma$ as a function of the effective bit-length $\ell_{\text{eff}} = \ell / n^{1/\omega}$ of the input RSA number. The data are consistent with an exponential decay, see Eq. (2), (dashed black line), characterized by fitted parameters $C_1 = 2.0 \pm 0.6$, $C_2 = 0.04 \pm 0.01$, and $\mu = 1.38 \pm 0.05$. The gray-shaded area around the dashed black line accounts for the uncertainty in the fitted parameters and corresponds to a 3σ confidence interval. Each data point is obtained by averaging ρ_{sr} over ten random RSA keys. Each ρ_{sr} is estimated using $n_{\text{CVP}} = 50$ lattices. (b) Scaling of the number of qubits, see Eq. (3), as a function of the bit-length ℓ and γ . The hyperparameter γ tunes the number of samples extracted from each CVP Hamiltonian spectrum to achieve AsrPL $\rho_{sr} = 1$. The extrapolated number of qubits n is required for factoring in a polynomial number of lattices within the $B_2 = 2n\ell$ smoothness bound. Vertical dotted lines denote three RSA challenges [35]: RSA-250 (829 bits), marking the largest RSA key factorized to date [10], RSA-1024 and RSA-2048.

C. Sieving via tree tensor networks

The core of the TNSS is the OPES algorithm introduced in Ref. [17]. This algorithm enables efficient sampling of bit-strings in the tail of a distribution from a TTN state, here representing a broad superposition of low-energy eigenstates of the CVP Hamiltonian in Eq. (1), see Appendix A 3. The key feature of the OPES algorithm is its ability to avoid resampling the same bit-string regardless of its probability within the TTN state. This allows us to extract bit-strings with very low probabilities while preventing multiple samples of those with high probabilities, then optimizing the sieving of sr-pairs.

The effectiveness of this approach is demonstrated in Fig. 3, which shows the results of the sieving procedure on a single lattice with $n = 32$ for a semiprime with $\ell = 70$. The TN sieving efficiency also depends on the bond dimension m (see Appendix A 2), the size of the indices of each tensor in the network [16,22]. We set $m = 8$ and sample approximately ℓ^3 bit-strings. The scatterplots in Figs. 3(a) and 3(b) show both sr-pairs (green dots) and no-sr pairs (gray dots) in terms of their distance from the target $\|\mathbf{t} - \mathbf{b}_j\|$, i.e., the square root of the bit-string energy (x axis), and the size of the pair j in bits ℓ_j (y axis), see Appendix A 3. All pairs accumulate in a low-energy region (lattice points close to the target) and




FIG. 3. Sampling smooth-relation pairs from a single lattice via tree tensor networks. (a, b) Lattice points sampled from a CVP Hamiltonian using the efficient TTN OPES method [17] with bond dimension $m = 8$. The input RSA key has a bit-length of $\ell = 70$ (dashed black horizontal line), and the CVP problem is characterized by $n = 32$ qubits. The number of samples is approximately 70^3 . No-sr pairs are denoted by gray dots in (a), while green dots represent sr-pairs in (b). The x axis indicates the square root of the energy of the extracted eigenstate, i.e., the Euclidean distance to the target $\|\mathbf{t} - \mathbf{b}_j\|$ of the pair j identified with the lattice point $\mathbf{b}_j$. The y axis represents the size of the sampled pair ℓ_j measured in bits. The measured average number of sr-pairs per lattice (AsrPL) on the chosen smoothness bound $\pi_2 = 2n\ell = 4480$ is $\rho_{sr} = 432$. For both scatterplots, we also illustrate the probability distribution of the pairs with respect to their distance from the target (top panel) and their bit-length (right panel). (c) The probability of sampling the first 400 pairs, ordered from most to least probable for both sr (green bars) and no-sr (gray bars). (d) The maximum values of the TTN bond dimension m as a function of the number of qubits n . The data show the sublinear scaling with the number of qubits $m(n) \sim An^\zeta$ with fitted parameters $A = 6.6 \pm 0.7$ and $\zeta = 0.42 \pm 0.02 \approx 2/5$ (dashed dark blue curve).

cluster around the input bit-length $\ell = 70$ [13,25], as shown by the probability distributions of green and gray dots in the panels on the right and at the top of each scatterplot.

In Fig. 3(c), we plot the probability of sampling a bit-string encoding one of the sieved pairs. For both sr and no-sr cases, we consider the first 432 extracted pairs with the largest probability. As clearly shown, the pairs with higher probabilities (gray bars) are those not useful for factoring. However, our sampling strategy excludes the possibility of measuring the same state more than once, thus efficiently sampling sr-pairs with lower probabilities (green bars). Notably, bond dimension $m = 8$ is sufficient to obtain $\rho_{sr} = 432$ independent sr-pairs despite their small probabilities.

Lastly, Fig. 3(d) displays the maximum bond dimension m used to represent the variationally optimized many-body quantum state, i.e., the solution to the CVP Hamiltonian in Eq. (1), as a function of the number of qubits n . The fit curve (dashed blue curve) indicates a sublinear growth of the bond dimension with n . These plots demonstrate the scalability of the TTN approach in solving the classical optimization task of searching for sr-pairs within the spectrum of CVP Hamiltonians that are characterized by long-range glassy interactions.

D. Tensor network Schnorr's sieving computational cost

We determine the scaling of the number of operations in the tensor network approach to Schnorr's sieving. To this end, we assume $\rho_{sr} = 1$, resulting in approximately $n_{\text{CVP}} \sim 2n\ell$ needed for the successful factorization of RSA keys with bit-length ℓ using n qubits.

The leading complexity in Schnorr's original factoring method arises from Babai's algorithm [2,15]. In TNSS, we perform this step for each CVP to obtain the approximate closest vector $\mathbf{b}^{cl}$ to the target. The lattice point $\mathbf{b}^{cl}$ is then used to define the Hamiltonian in Eq. (1). The number of operations $\mathcal{T}_1$ to execute Babai's algorithm scales as

$$\mathcal{T}_1 = \mathcal{O}(n^7 \ell \log_{10}^3 n). \quad (4)$$

The TN sieving step involves the variational energy optimization of the initial TTN state with bond dimension m , followed by the OPES sampling of lattice points for each CVP. If we sample ℓ^γ pairs per CVP, the TN-based part $\mathcal{T}_2$ accounts for the following number of operations (see Appendix A 3 for details)

$$\mathcal{T}_2 = \mathcal{O}(n^4 \ell m^4 + n^2 \ell^{\gamma+1} m^4). \quad (5)$$

Once the pairs are sampled, we construct the corresponding integers by computing the lattice point components on the basis $\mathcal{B}$ (see Appendix A 1). The sr-pairs are identified by decomposing these integers on the smoothness basis P_2 . The worst-case scenario corresponds to applying a trial division strategy. In this case, the total complexity $\mathcal{T}_3$ is bounded by

$$\mathcal{T}_3 = \mathcal{O}(n^4 \ell + n^3 \ell^{\gamma+1} + n^2 \ell^{\gamma+2}). \quad (6)$$

Given that the bond dimension m grows as $n^{2/5}$, according to Fig. 3(d), the time complexity $\mathcal{T}_3$ in Eq. (6) is always subleading compared to the TN sieving step [see scaling in Eq. (5)]. Overall, the number of operations $\mathcal{T} = \mathcal{T}_1 + \mathcal{T}_2$ scales as

$$\mathcal{T} = \mathcal{O}(n^7 \ell \log_{10}^3 n + n^4 \ell m^4 + n^2 \ell^{\gamma+1} m^4). \quad (7)$$




FIG. 4. Scaling of the number of operations for the TNSS algorithm. Scaling of the number of operations as a function of the input size ℓ for the TNSS algorithm derived from the predicted scaling of the number of qubits in Eqs. (3) and (7). We set the average number of sr-pairs per lattice (AsrPL) in Eq. (3) to $\rho_{sr} = 1$. Solid curves represent Babai's contribution $\mathcal{T}_1$ [first term in Eq. (7)], dashed curves the TTN sieving term $\mathcal{T}_2$ in Eq. (7), for various values of the hyperparameter γ . Vertical dotted lines denote three RSA challenges [35]: RSA-250 (829 bits) marking the largest RSA key factorized to date [10], RSA-1024, and RSA-2048. The leading contribution in the algorithm varies with the number of sampled pairs: below $\gamma = 8$, Babai's algorithm prevails over the computation required for sampling lattice points in the range of currently relevant bit-lengths of $\ell \sim 10^3$. Within this range, the TN term $\mathcal{T}_2$ in Eq. (7) becomes the leading one for $\gamma > 8$.

The balance between Babai's complexity $\mathcal{T}_1$ and TTN complexity $\mathcal{T}_2$ in Eq. (7) depends on the input bit-length ℓ and the number of eigenstates sampled to obtain a sufficient set of sr-pairs, as determined by the hyperparameter γ . As illustrated in Fig. 4, Babai's algorithm dominates up to $\gamma < 8$ for relevant key lengths of approximately $\ell \sim 10^3$.

In conclusion, we achieve an asymptotic polynomial scaling of the number of operations $\mathcal{T} \lesssim \mathcal{O}(\ell^{59} \log_{10}^3 \ell)$ for $\gamma < 25$ based on the scaling for the number of qubits in Eq. (3). Furthermore, the memory requirements for the TN sieving process remain polynomial. The scaling of the bond dimension m with the number of qubits shown in Fig. 3(d) ensures that the wave-function description of CVP solutions does not scale exponentially with n . For a single CVP, the memory needed for the TTN state and the Hamiltonian operator scales at most as $\mathcal{O}(nm^3) + \mathcal{O}(n^3 m^2) \lesssim \mathcal{O}(\ell^{32})$ bits, while the sampling routine stores up to $\mathcal{O}(\ell^\gamma)$ n -bit strings.

When compared to the GNFS [8,37], the TNSS scaling of the number of operations is larger within the RSA key bit-length range relevant to current security protocols. Consequently, the GNFS remains the leading classical algorithm for factoring RSA. However, in the limit of large ℓ , the TNSS approach becomes more efficient due to its predicted polynomial complexity.

E. Conclusion

We introduce a classical algorithm for RSA key factorization based on tensor network sieving applied to Schnorr's lattice method. We achieve successful factorization by

fine-tuning the hyperparameters of the TNSS algorithm, e.g., the number of qubits n , the smoothness basis size π_2 , and the number of extracted eigenstates via γ . The algorithm effectiveness spawns from the efficient TTN optimization and the OPES sampling.

The TNSS algorithm successfully factorizes randomly generated RSA numbers up to 100 bits, validating the full pipeline of lattice generation, sieving, and smooth-related pairs postprocessing. To the best of our knowledge, this is the largest RSA number factorized using Schnorr’s sieving method.

The algorithm is quantum inspired as it leverages concepts and methods from quantum algorithms. However, due to the tensor network approach and the limited amount of quantum correlations generated during the process, it can be efficiently run on classical computers. We foresee that employing a universal scalable quantum computer to execute the same algorithm would not be beneficial: due to the distribution shown in Fig. 3, the sampling of the classical configurations on the quantum computer’s wave function would be highly inefficient, necessitating the development and exploration of alternative strategies.

We emphasize that the TNSS algorithm’s limitation in factoring larger keys lies in the high-order polynomial scaling of resources. We identify Babai’s algorithm and the TN sampling of sr-pairs as the primary contributors to the overall scaling. However, TN efficient OPES sampling has emerged as crucial for acquiring sr-pairs useful for factoring N . This is mainly due to its ability to avoid resampling.

The numerical simulations are executed on a single-node machine and will benefit from massive parallelization on multinode architectures. Indeed, the algorithm features several highly parallelizable steps, such as the TN sampling and simultaneous CVP sieving: leveraging high-performance computing, parallelization, and code optimization could significantly improve its performance.

In conclusion, we numerically investigate the scaling of classical resources required by the TNSS algorithm and observed behavior consistent with polynomial scaling over the studied range of bit-lengths $\ell \in [50, 130]$. These findings suggest that tensor networks could provide a novel computational framework for addressing lattice-based cryptographic problems. While our conclusions are based on numerical evidence, and not on formal complexity-theoretical proofs, the observed trends indicate that factoring RSA integers using this approach may be achievable with polynomial classical resources, pending further validation at larger input sizes or through analytical investigation.

ACKNOWLEDGMENTS

We thank M. Ballarin, M. Di Liberto, M. Rigobello, P. Silvi, and E. Altelaarrea-Ferré for useful discussions. We thank R. Fazio, T. Osborne, M. Montangero, S. Pascasio, F. Pollmann, N. Schuch, M. Stoudenmire, F. Verstraete, and P. Zoller for feedback on our work and on the manuscript. We thank L. Zanfrando for technical support on the computational side.

This work is partially supported by the World Class Research Infrastructure-Quantum Computing and Simulation Center (QCSC) of Padova University, Quantum Frontiers

and QuaSiModO (the Departments of Excellence 2023-2027) the Italian Ministry of University and Research (MUR) via project TANQU (PRIN2022), the Italian Istituto Nazionale di Fisica Nucleare (INFN) via project Iniziativa Specifica IS-Quantum, the QuantERA projects TNISQ and QuantHEP, the European Union Horizon 2020 research and innovation program (Quantum Flagship PASQuanS2, Euryqa), the Italian Research Center on HPC, Big Data and Quantum Computing ICSC (NextGenerationEU Project No. CN00000013), the program PON “Ricerca e Innovazione” 2014-2020 Project No. 19-G-12542-2, the University of Bari through the 2023-UNBACLE-0244025 grant, the German Federal Ministry of Education and Research (BMBF) via the project QRydDemo, and the project 101194322 – QEC4QEA – HORIZON-EUROHPC-JU-2023-QEC-05. We acknowledge computational resources from Cloud Veneto, CINECA, BwUniCluster, and the 2017 Strategic Research Infrastructure Grant “CAPRI Calcolo ad Alte Prestazioni per la Ricerca e l’Innovazione” of Padova University.

DATA AVAILABILITY

The data that support the findings of this article are available upon reasonable request from the authors.

APPENDIX

The appendix includes a detailed introduction to Schnorr’s lattice sieving in Appendix A1, the description of our tensor network formulation in Appendix A2 and of the sampling formalism in OPES in Appendix A3. In Appendix A4, we validate and compare our approach to the original Schnorr’s formulations and Yan *et al.* extension [15]. In Appendix A5 we provide the prime factors of our record RSA-key factorization with the TNSS algorithm. Finally, in Appendix A6 we describe the fitting procedure and how to choose the functions to rescale the bit-lengths in Fig. 2(a).

1. Schnorr’s lattice-based factoring algorithm

In this appendix, we describe in detail Schnorr’s original procedure [13,18] for decomposing large semiprimes, as well as the quantum variant proposed in Ref. [15].

Factoring by congruence of squares. Schnorr’s factoring algorithm is rooted in the congruence of squares method [2,4]. Given a semiprime number $N = pq$, the algorithm determines the prime factors p and q by constructing the congruence $X^2 \equiv Y^2 \pmod{N}$, with $X \not\equiv \pm Y \pmod{N}$ being integers. Then, the algorithm efficiently computes $p = \gcd(X + Y, N)$ and $q = \gcd(X - Y, N)$ using the Euclidean algorithm [38]. Many classical and quantum factorization techniques [6,7,33,37], including Shor’s algorithm [11], employ this strategy, differing in how they construct the above congruence. In Schnorr’s integer factorization, this process comprises two stages: an initial *collection phase* followed by subsequent *processing phase*.

Collection phase: Smooth numbers. The objective of the collection phase is to accumulate a sufficient number of so-called *smooth-relation pairs* (sr-pairs). These pairs are defined based on smooth numbers, a pivotal concept in number theory and essential for factoring tasks [4,7]. We introduce the set of

the first π consecutive prime numbers

$$P = \{p_1, p_2, \dots, p_\pi\} \quad (\text{A1})$$

and call $B = p_\pi$ the smoothness bound. An integer is said to be smooth with respect to this bound if all its prime factors are less than or equal to p_π . Every B -smooth integer u can be represented on the prime basis P as

$$u = \prod_{p_j \in P} p_j^{e_j}. \quad (\text{A2})$$

The exponents $\{e_j\}_{j=1}^\pi \in \mathbb{Z}^\pi$ define the multiplicities in the prime decomposition of u . Negative B -smooth integers are equivalently represented on the prime basis P by including an extra element $p_0 = -1$ in Eq. (A1). Thus, negative numbers are distinguished by the exponent $e_0 = 1$, while positive numbers have $e_0 = 0$. As the prime decomposition is unique, we can represent a smooth number by the corresponding multiplicity vector $u \sim \mathbf{e} = (e_0, e_1, \dots, e_\pi)$. We define an integer pair (u, v) as a *smooth pair* if both u and v are B smooth. If the integer $u - vN$ also turns out to be B smooth, the triplet $(u, v, u - vN)$ is mathematically referred to as a *fac-relation* [18]. We denote $(u, u - vN)$ as a smooth-relation (sr) pair. By construction, the integers forming an sr-pair satisfy $u \equiv u - vN \pmod{N}$, leading to the following congruence:

$$\frac{u - vN}{u} \equiv 1 \pmod{N}. \quad (\text{A3})$$

Additionally, we can represent the ratio of the B -smooth integers $u - vN$ and u in Eq. (A3) using the corresponding multiplicity vectors on the common prime basis P

$$\begin{aligned} \frac{u - vN}{u} &= \frac{\prod_{p_j \in P} p_j^{e'_j}}{\prod_{p_j \in P} p_j^{e_j}} \\ &= \prod_{p_j \in P} p_j^{(e'_j - e_j)} = \prod_{p_j \in P} p_j^{\tilde{e}_j} \\ &\sim \tilde{\mathbf{e}} = (\tilde{e}_0, \tilde{e}_1, \dots, \tilde{e}_\pi). \end{aligned} \quad (\text{A4})$$

The most computationally demanding part of the factoring procedure consists of collecting a sufficient number D of sr-pairs $\{(u, u - vN)\}_{r=1}^D$, where D depends on the smoothness bound B [33].

Processing phase via linear algebra in $\mathbb{GF}(2)$. Schnorr's approach involves leveraging the simple multiplicative structure of sr-pairs $(u - vN) / u$ to construct two congruent squares X^2 and Y^2 , thereby factorizing N . Once D ratios of the form defined in Eq. (A4) are collected, the algorithm quests for a combination that results in a ratio between two square integers

$$\prod_{r=1}^D \tau_r \left(\frac{u_r - v_r N}{u_r} \right) = \prod_{p_j \in P} p_j^{\sum_{r=1}^D \tau_r \tilde{e}_{j,r}} = \left(\frac{X}{Y} \right)^2. \quad (\text{A5})$$

Here, $\tau \in \{0, 1\}^D$ is a binary vector to be determined. The product of B -smooth numbers on the left-hand side of Eq. (A5) generates another B -smooth number. Given the property in Eq. (A3), applying the modulo N operation on both sides of Eq. (A5) produces the desired congruence $X^2 / Y^2 \equiv 1 \pmod{N}$. To obtain a combination of smooth numbers with only even multiplicities, i.e., a square integer, we need to determine the binary vector τ . This is accomplished by solving

the following system of linear equations [2,4,6]:

$$\sum_{r=1}^D \tau_r \tilde{e}_{j,r} \equiv 0 \pmod{2}, \quad (\text{A6})$$

for $j = 0, \dots, \pi$. The system can be represented in matrix form as

$$\mathcal{E} \tau = 0, \quad (\text{A7})$$

transforming the problem into finding the kernel of a binary $(\pi + 1) \times D$ matrix $\mathcal{E}$, where the r th column corresponds to the multiplicity vector $\tilde{\mathbf{e}}$ of the r th collected sr-pair modulo 2. Theorems in linear algebra on the Galois field $\mathbb{GF}(2)$ of order 2 ensure that a nontrivial solution to Eq. (A7) exists when the system has at least one free variable, i.e., when $D = \pi + 2$ [2,32]. This defines the minimum number of sr-pairs to be collected during the first stage of the algorithm. This postcollection processing step is a minor but common part of many integer factoring methods [5,7,8,33,37] and efficient techniques are available to find the solution vector τ [9,10,26]. It can be rigorously demonstrated [4,33] that at least half of the solutions τ to Eq. (A7) which yield the congruence $X^2 \equiv Y^2 \pmod{N}$ also meet the nontriviality condition $X \not\equiv \pm Y \pmod{N}$, thus leading to the factorization of N .

Collecting sr-pairs through lattice sieving. Within this framework, Schnorr first proposed a lattice-based algorithm in Ref. [18], where the task of collecting at least $D = \pi + 2$ sr-pairs is reframed as solving a series of optimization problems known as closest vector problems (CVPs). A CVP entails determining the lattice point $\mathbf{b}^{op}$ in a given lattice Λ that is closest to a specified target point $\mathbf{t}$ [2]. Mathematically, we define the π -rank $(\pi + 1)$ -dimensional lattice as the following infinite set of discrete points:

$$\Lambda(\mathcal{B}) = \left\{ \sum_{j=1}^\pi e_j \mathbf{b}_j \mid \{e_j\}_{j=1}^\pi \in \mathbb{Z}^\pi \right\}, \quad (\text{A8})$$

generated by the lattice basis

$$\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_\pi\}. \quad (\text{A9})$$

Each $\mathbf{b}_j \in \mathbb{Z}^{\pi+1}$ for $j = 1, 2, \dots, \pi$, and $\mathbf{t} \in \mathbb{Z}^{\pi+1}$ is a target point not in the lattice $\Lambda(\mathcal{B})$. The problem asks for the lattice point $\mathbf{b}^{op} \in \Lambda(\mathcal{B})$ satisfying

$$\|\mathbf{b}^{op} - \mathbf{t}\| = \min_{\mathbf{b} \in \Lambda(\mathcal{B})} \|\mathbf{b} - \mathbf{t}\|, \quad (\text{A10})$$

in the Euclidean ℓ_2 norm. Solving a generic instance of the CVP is NP hard [19,20]. However, approximate techniques obtain nearly optimal solutions. For instance, Babai's nearest plane algorithm [2,14] solves this problem in polynomial time with the lattice rank π , provided Eq. (A10) is relaxed with an approximation factor $\eta = 2^{\mathcal{O}(\pi)}$. The algorithm finds an approximate closest vector $\mathbf{b}^{cl}$ such that

$$\|\mathbf{b}^{cl} - \mathbf{t}\| \leq \eta \times \min_{\mathbf{b} \in \Lambda(\mathcal{B})} \|\mathbf{b} - \mathbf{t}\|, \quad (\text{A11})$$

within a number of operations that scales as shown in Ref. [2] with

$$\mathcal{O}(\pi^6 [\log_{10} (\max_{\mathbf{b}_j \in \mathcal{B}} \|\mathbf{b}_j\|)]^3). \quad (\text{A12})$$

Schnorr's factoring algorithm [13,15,18] gets independent sr-pairs from approximate closest vector solutions to the N -related target

$$\mathbf{t}_{N,c} = (0 \quad \cdots \quad 0 \quad \lceil 10^c \ln N \rceil)^T, \quad (\text{A13})$$

on a set of lattices described by the following $(\pi + 1) \times \pi$ lattice basis matrices:

$$\mathcal{B}_{f,c} = \begin{pmatrix} f(1) & 0 & \cdots & 0 \\ 0 & f(2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & f(\pi) \\ \lceil 10^c \ln p_1 \rceil & \lceil 10^c \ln p_2 \rceil & \cdots & \lceil 10^c \ln p_\pi \rceil \end{pmatrix}, \quad (\text{A14})$$

where p_j for $j = 1, \dots, \pi$ are the prime numbers in the factoring basis P . The columns of the lattice basis matrix in Eq. (A14) are built off the lattice basis vectors $\mathbf{b}_j \in \mathbb{Z}^{\pi+1}$. A set of parameters can be varied to generate different CVPs. Specifically, c is the so-called lattice precision parameter, while different diagonals of $\mathcal{B}$ are generated via a random permutation of the set $\{\lceil j/2 \rceil\}_{j=1}^\pi, f(j)$ being the j th element in the permuted set. The symbol $\lceil \cdot \rceil$ denotes the nearest integer rounding operation, with the convention that numbers with a fractional part of 0.5 are rounded upwards. In this lattice mapping of the factoring problem, each lattice point $\mathbf{b} \in \Lambda(\mathcal{B})$ corresponds uniquely to a pair of p_π -smooth integers (u, v) . Moreover, those lattice points closest to the target in each CVP $(\Lambda, \mathbf{t})$ potentially offer sr-pairs. This stems from the above lattice construction, where sr-pairs represent approximate solutions to a system of simultaneous linear Diophantine equations for $\ln N$ [18]. This factoring technique is known as *Schnorr's lattice sieving* because it filters sr-pairs that satisfy specific linear Diophantine equations for $\ln N$ using lattice-target pairs $(\Lambda, \mathbf{t})$ associated with N .

Constructing sr-pairs from lattice points. B -smooth pairs (u, v) of positive integers are obtained from any lattice point $\mathbf{b} = \sum_{j=1}^\pi e_j \mathbf{b}_j$ as follows:

$$u = \prod_{\substack{e_j \geq 0 \\ j \in [1, \pi]}} p_j^{e_j}, \quad v = \prod_{\substack{e_j < 0 \\ j \in [1, \pi]}} p_j^{-e_j}. \quad (\text{A15})$$

Hence, the component vector $\mathbf{e} \in \mathbb{Z}^\pi$ of $\mathbf{b}$ on the lattice basis $\mathcal{B}$ translates to the multiplicity vectors of the integers u and v . Given a lattice point, we compute these multiplicities by inverting the lattice basis matrix

$$\mathbf{e} = \mathcal{B}^{-1} \mathbf{b}. \quad (\text{A16})$$

This calculation is essential in Schnorr's lattice sieving. The number of operations required for computing the matrix pseudoinverse scales as $\mathcal{O}(\pi^3)$ for each lattice. Additionally, $\mathcal{O}(\pi^2)$ operations are needed for the matrix-vector multiplication in Eq. (A16) each time we construct a smooth pair from a lattice point. From (u, v) , we form the integer $u - vN$ and decompose it into its prime factors to determine whether it is B smooth, i.e., whether the lattice point $\mathbf{b} \sim (u - vN) / u \sim \tilde{\mathbf{e}}$ yields an sr-pair for factoring. The time complexity for each $u - vN$ prime decomposition is bounded by a trial division strategy on the smoothness basis of size π , i.e., $\mathcal{O}(\pi)$,

assuming finding remainder and division take place in constant time [4].

Schnorr's sublinear theory. Drawing from number theory and heuristic arguments, Schnorr proposed constructing CVPs with sublinear lattice dimensions [13,18]

$$\pi = \frac{\ell}{\log_2 \ell}, \quad (\text{A17})$$

alongside the corresponding sublinear smoothness bound for factoring N . Here $\ell = \lfloor \log_2(N + 1) \rfloor$ represents the size in bits of the input semiprime N and $\lfloor \cdot \rfloor$ denotes the floor function. Schnorr asserted that this sublinear bound, combined with efficient lattice reduction methods to compute approximate CVP solutions, would break the security of RSA-2048 more efficiently than state-of-the-art classical methods [13,18]. However, this theoretical statement has sparked significant debate. Practical implementations of Schnorr's sublinear classical algorithm have revealed discrepancies with his prediction [34]. The primary shortcoming arises from the probability of a CVP approximate solution being associated with an sr-pair, which decays exponentially as the input bit-length increases.

Sublinear-resource quantum-classical Schnorr's sieving. To address the exponential complexity of the classical search for sr-pairs, Yan *et al.* recently proposed in Ref. [15] a hybrid quantum-classical variant of Schnorr's approach. This *sublinear-resource quantum integer factorization* method combines classical Babai's lattice reduction with quantum optimization, specifically the quantum approximate optimization algorithm (QAOA) [39,40]. The goal is to extract more sr-pairs from each CVP by considering not only Babai's solution $\mathbf{b}^{cl}$ but also a collection of "quantum" solutions that offer additional approximate closest lattice points to the target. The quantum approach builds on the following cost function:

$$F(x_1, \dots, x_\pi) = \left\| \mathbf{t} - \mathbf{b}^{cl} - \sum_{j=1}^\pi \kappa_j \mathbf{d}_j \right\|^2, \quad (\text{A18})$$

where $\boldsymbol{\kappa} = (\kappa_1, \dots, \kappa_\pi) \in \{0, \pm 1\}^\pi$ are rounding optimization variables and $\mathcal{D} = \{\mathbf{d}_1, \dots, \mathbf{d}_\pi\}$ is the LLL-reduced lattice basis of $\mathcal{B}$ obtained through the LLL algorithm [2,28]. Babai's method is based on the LLL lattice reduction and outputs the lattice point $\mathbf{b}^{cl} = \sum_{j=1}^\pi c_j \mathbf{d}_j$, with $c_j = \lfloor \mu_j \rfloor$ given by the nearest integer to the Gram-Schmidt coefficient obtained from the orthogonalization of $\mathcal{B}$ [14]. Once the cost function in Eq. (A18) is minimized, the rounding variables provide a more accurate or at least equivalent closest vector to the target

$$\mathbf{b}^q = \mathbf{b}^{cl} + \sum_{j=1}^\pi \kappa_j \mathbf{d}_j = \sum_{j=1}^\pi (c_j + \kappa_j) \mathbf{d}_j. \quad (\text{A19})$$

To encode efficiently the possible values $-1, 0$ and $+1$, i.e., all the possible rounding corrections to $\mathbf{b}^{cl}$ in Eq. (A19), each $\kappa_j \in \{0, \pm 1\}$ is promoted to a quantum spin-1/2 operator, i.e., qubits,

$$\kappa_j \rightarrow \hat{\kappa}_j = \frac{1}{2} \text{sign}(\mu_j - c_j) (\mathbb{1} - \hat{\sigma}_j^z), \quad (\text{A20})$$

where $\hat{\sigma}_j^z$ denotes the z -Pauli matrix acting on the j th qubit. We obtain the spin-glass Hamiltonian encoding the

optimization problem defined in Eq. (A18)

$$\mathcal{H} = \sum_{k=1}^{\pi+1} \left(t_k - b_k^{cl} - \sum_{j=1}^{\pi} d_{j,k} \kappa_j \right)^2, \quad (\text{A21})$$

which describes a system of $n = \pi$ all-to-all interacting qubits. For simplicity, we omit the quantum operator symbol whenever it is clear from the context. The Hamiltonian in Eq. (A21) takes in all potential roundings of the Gram-Schmidt coefficients within Babai's nearest plane algorithm. In a scenario where computing the entire spectrum of an n -qubit system is feasible, all 2^n possible roundings around $\mathbf{b}^{cl}$ are available. These new lattice points can be used to construct smooth pairs and potentially collect more than one sr-pairs per CVP. Note that this approach involves accessing the excited states of $\mathcal{H}$ for each CVP, rather than only the ground-state encoding the minimum of the cost function, as typically done in quantum approaches to classical optimization [41].

2. Tensor network formulation

TNs are powerful variational ansätze tailored to efficiently represent quantum states taking advantage of the correlations and entanglement properties that characterize these states [16,21,22,42]. Numerical methods rooted in TNs have demonstrated remarkable success in classically simulating the equilibrium and out-of-equilibrium behavior of quantum many-body systems, spanning from condensed matter to lattice gauge theories [43–47]. TN methods thus provide a representation of the exponentially large wave function of a quantum many-body system using the Schmidt decomposition to compress the entanglement content of the wave function. The efficiency of this decomposition is governed by the so-called bond dimension m , i.e., the Schmidt rank, which interpolates between a product state with $m = 1$ and the exact, albeit unfeasible, state with $m = 2^{n/2}$.

Among the various TN geometries available, our paper focuses on *tree tensor networks* (TTNs). TTNs are among the most prominent geometries for investigating long-range interacting many-body systems due to their hierarchical structure, improved connectivity, and computational efficiency regarding both bond dimension and system size [16]. TTNs have recently shown promise in addressing hard classical optimization problems [48]. These problems are typically associated with classical Hamiltonians that exhibit long-range two-body interactions among Ising spins [41]. In the context of our research, we aim to leverage the TTN adaptability to enhance the demanding optimization task related to the search for sr-pairs within the spectrum of the CVP Hamiltonian $\mathcal{H}$ defined in Eq. (A21), which is a classical Hamiltonian, being diagonal within the chosen computational basis, and characterized by glassy all-to-all interactions.

Typically, the variational ground-state search (used in our approach as well) approximates the ground state of a physical system with a specific ansatz, e.g., a TTN state. As optimizations are local with respect to the tensor in the network and the bond dimension limits the entanglement, it usually takes many sweeps to converge to the ground state. Replicating this approach without modifications, we sample exactly one eigenstate of the system; the diagonal Hamiltonian has

eigenstates with $m = 1$ which converges well. Established methods like the excited state search do either scale unfavorably when requesting many excited states [49] or we have not explored them like the minimal entangled typical thermal state algorithm [50]. Appendix A 3 explains the sampling method that we use for the tensor-network-enhanced Schnorr's algorithm presented here; the algorithm fulfills the requirements in terms of the number of sampled states required for factoring.

The detailed scaling of the variational ground-state search is

$$\mathcal{O}(n^3 m^4), \quad (\text{A22})$$

where the contraction of an effective operator with a TTN tensor is the most costly step with $\mathcal{O}(m^4)$; one power of n originates in the number of tensors to be updated, the other two powers from the bond dimension of the Hamiltonian representation. We ignore the number of sweeps due to the quick convergence.

3. Sampling formalism

The sampling formalism relies on three key aspects of the problem itself and tensor network methods: (i) the eigenstates are also eigenstates of the z basis, i.e., any projective measurement in the z basis yields a potential solution; (ii) an insufficiently converged TTN ground state overlaps with the excited states, although the probabilities are small; (iii) we can efficiently sample states in decaying probability tails with the algorithm presented in Ref. [17]. Thus, our goal is to launch a variational ground-state search which is not well converged; the sampling yields then excited states in the z basis which are potentially low-energy smooth-relation pairs for Schnorr's algorithm.

As the TTN converges well with the diagonal Hamiltonian $\mathcal{H}$ even within one sweep, we introduce a perturbation breaking the diagonal form via random transverse fields $\{h_x(j)\}_{j=1}^n$,

$$\mathcal{H}' = \mathcal{H} + \sum_{j=1}^n h_x(j) \sigma_j^x. \quad (\text{A23})$$

We use random couplings $h_x(j)$ for the Pauli operator σ_j^x acting on site j . Specifically, the strengths of the local random transverse fields are controlled by a tunable hyperparameter $\alpha \in \mathbb{R}_+$, defined as

$$h_x(j) \sim \mathcal{U}(-|\bar{h}_x|, |\bar{h}_x|) \\ \bar{h}_x \sim \frac{1}{\alpha} \{ \{W_{jj'}\}, \{h_j\} \}. \quad (\text{A24})$$

The parameter $\alpha \in \mathbb{R}_+$ sets the relative magnitude of the transverse fields with respect to the diagonal spin-glass couplings $\{W_{jj'}\}_{j,j'=1}^n$ and $\{h_j\}_{j=1}^n$ —corresponding to the two-body spin interactions and single-spin bias terms in Eq. (A21), respectively. Here, $\mathcal{U}$ denotes the uniform distribution over the indicated interval, while $\langle \circ \rangle$ is the arithmetic mean. By adjusting α , we control the scale of the off-diagonal (transverse) terms relative to the diagonal couplings. Computationally, we define the special case $\alpha = 0$ as the one corresponding to omitting the transverse-field term entirely, recovering the classical unperturbed CVP Hamiltonian in Eq. (A21). We carefully tune the amplitude of the random couplings $h_x(j)$




FIG. 5. Relation between bond dimension and sampling smooth-relation pairs from a single lattice via tree tensor networks. We plot both no-sr (gray points) and sr-pairs (green points) sampled from a lattice using the efficient TTN OPES method at different bond dimensions m in the first row. The input RSA key has a bit-length of $\ell = 70$, and the CVP problem is characterized by $n = 32$. The number of samples is approximately $\ell^3 = 70^3$. The x axis indicates the square root of the energy $\|t - b_j\|$ of the extracted eigenstate, i.e., the Euclidean distance to target t of the sampled pair identified with the lattice point b_j . The y axis represents the bit-length ℓ_j of the integer $u - vN$ associated with the sampled pair. Bond dimensions greater than $m = 8$ do not notably impact the sieving performance. Indeed, the average number of sr-pairs per lattice (AsrPL), i.e., the number of green points, is approximately $\rho_{sr} \approx 400$ for each bond dimension. The second row shows the corresponding sampling probability for the first 400 no-sr and sr-pairs ordered by decreasing probability. The x axis denotes the pair's numbering index j , while the y axis shows the probability of the associated bit-string (gray bars for no-sr pairs, green bars for sr-pairs).

and then obtain TTN wave functions with sufficient overlap to excited states of the original Hamiltonian.

We use the sampling algorithm from Ref. [17] out-of-the-box. Multiple superiterations are carried out with exit criteria for the number of target states or the accumulated probability from sampled states. The ability to avoid resampling known states leads to an exponential speedup for sampling small probabilities in exponentially decaying tails of distributions. We present the dataset of samples and their probability histogram at different bond dimensions m in Fig. 5, where Figs. 3(a) to 3(c) in Sec. IC contain the subset of this data for the $m = 8$ case. Recall that the data are for our sieving procedure on a single lattice with $\ell = 70$ and $n = 32$. We plot the quality of the sampled pairs in terms of their distance from the target point (x axis) and the size of the resulting integer $u - vN$ in bits. We sample approximately ℓ^3 states. We conclude that the impact of the bond dimension is minor here, while this parameter usually is key to convergence in simulations of entangled systems, e.g., condensed matter problems or lattice gauge theories. Note that a bond dimension $m = 8$ is sufficient to obtain an AsrPL of approximately $\rho_{sr} = 400$ independent sr-pairs from a single lattice, even though their probabilities of being measured are on the order of 10^{-6} .

The detailed scaling for the sampling including the ground-state search of Eq. (A22) yields

$$\mathcal{O}(n^3 m^4 + n \ell^\gamma m^4), \quad (\text{A25})$$

where we directly enter the number of samples as the bit-length ℓ and the meta parameter γ . The sampling of the TTN state is independent of the bond dimension of the Hamiltonian: hence, the cost scales linearly with the number of qubits.

In the future, other techniques than the ground-state search can be exploited in combination with the sampling, e.g., finite-temperature representations [51–53]. The sampling of states can also extend beyond tensor network methods. Notably, this approach is also valuable for benchmarking quantum computer algorithms or quantum simulators devoted to Schnorr's sieving. In particular, it is useful to understand the limitations and bottlenecks to be considered and overcome.

4. Beyond Schnorr's sublinear theory

In this appendix, we numerically demonstrate the potential advantage of the hybrid classical-quantum lattice sieving outlined in Ref. [15] over the original Schnorr's factoring algorithm, provided that the sublinear scaling stated in




FIG. 6. Comparison between Schnorr's classical sieving and the quantum-powered sieving. We compare the resource requirement of different algorithms: the classical Schnorr's algorithm (Babai) [13] with lattice size and smoothness bound beyond the original sublinear assumption (black squares); Yan *et al.* [15] implementation of the quantum version of Schnorr's algorithm (blue diamonds) with sublinear scaling of the lattice size and smoothness bound; Yan *et al.* [15] implementation of the quantum version of Schnorr's algorithm (red triangles) beyond sublinear scaling of the lattice and smoothness bound. The x axes represent the bit-length of the input RSA number, while the y axes carry the chosen relevant quantities to the considered algorithms, i.e., the number of CVPs n_{CVP} , left panel, and the average number of sr-pairs per lattice ρ_{sr} , right panel. The implementation of the classical Schnorr's algorithm (Babai) with sublinear resources does not appear in the plots as the factorization is never successful.

Ref. [13]

$$\pi_1 = \pi_2 = \frac{\ell}{\log_2 \ell} \quad (\text{A26})$$

is relaxed. First, we implement Schnorr's sublinear method to factorize RSA keys with increasing bit-lengths in the range $\ell \in [10, 60]$. Using the size defined in Eq. (A26) for both the factoring and the smoothness bases, the algorithm does not successfully factorize any of the considered bit-lengths. This is due to an insufficient number of sr-pairs found while sieving Babai's solutions through sublinear dimensional lattices and evaluating their smoothness on $B = p_{\pi_2}$.

To increase the number of sr-pairs, we extend beyond the sublinear theory by setting the lattice rank π_1 to the theoretical sublinear value and gradually increase it by 1 until a nontrivial factorization is obtained. The size of the smoothness basis is adjusted to $\pi_2 = 2\pi_1^2$ as arbitrarily proposed in Ref. [15], albeit without any justification. In Fig. 6, we plot the AsrPL and the number of CVPs necessary to compute the two prime factors, p and q , of the RSA keys with bit-lengths up to $\ell = 60$. The values of the increased lattice rank π_1 used in these simulations are summarized in Table I. The black squares indicate an exponentially vanishing ρ_{sr} , resulting in an exponential number of CVPs required for factoring. This demonstrates that even relaxing the sublinear growth of π_1 and the smoothness bound B with problem size, Schnorr's lattice sieving requires an exponential number of resources.

Following the authors of Ref. [15], we reexamine Schnorr's sublinear theory by implementing the spin glass mapping of the CVP set to factorize the same RSA keys. In this method, the lattice rank π_1 corresponds to the number of qubits n . Given the small values of $n = \ell / \log_2 \ell$ when considering $\ell \in [10, 60]$, we exactly compute the 2^n spin-glass

configurations in the spectrum of each CVP Hamiltonian defined in Eq. (A21), and thus obtain the AsrPL and the number of CVPs Hamiltonian as before. The blue diamonds in Fig. 6 show an improvement compared to the previous Schnorr's factorizations. Despite achieving successful factorization with a sublinear number of qubits, the exponential decay (growth) of ρ_{sr} (n_{CVP}) persists. This issue has been investigated in other works, see Refs. [23–25].

Finally, we increase the sublinear bounds again following Table I to provide a fairer comparison between the spin-glass-based sieving and Schnorr's algorithm. The red triangles in Fig. 6 suggest that by increasing the number of qubits beyond the sublinear scaling, we can shift the exponential decay of ρ_{sr} to larger bit-lengths. Indeed, both ρ_{sr} and n_{CVP} appear constant in the plotted range of ℓ . To obtain these data points, we access the full CVP spectra as before.

TABLE I. Lattice ranks, i.e., the number of qubits, used to obtain the results in Fig. 6. When the sublinear theory defined in Eq. (A26) is not enough to successfully factorize an RSA key N with bit-length ℓ , we incrementally increase π_1 (n in the classical-quantum Schnorr's sieving) and set $\pi_2 = 2\pi_1^2$ until nontrivial prime factors p and q of N are obtained.

Bit-length ℓ	Sublinear lattice rank π_1	Increased lattice rank π_1
10	3	4
20	5	7
30	6	8
40	8	10
50	9	12
60	10	15

This preliminary investigation of the two lattice sieving methods prompted us to introduce the TNSS and use it to scale the simulations with n , thereby characterizing the computational resources needed to target RSA keys of significant size for modern cryptosystems.

5. 100-bit factorization

Here, we report the details regarding the largest factorization of the 100-bit RSA key

$$N = 791\,339\,171\,587\,617\,359\,026\,543\,582\,309 \quad (\text{A27})$$

achieved using the TNSS algorithm. We sieve $n_{\text{CVP}} = 28\,980$ CVPs characterized by $n = 64$. We set $\gamma = 2$ and find 11 278 sr-pairs on the smoothness bound $B = 137\,477$ using a basis of prime numbers of size $\pi_2 = 12\,801$. By combining these sr-pairs, we obtain the following factorization:

$$N = 428\,949\,705\,601\,033 \times 1\,844\,829\,734\,709\,373. \quad (\text{A28})$$

Notice that the number of sr-pairs needed to find a nontrivial solution is less than the theoretical value given by $\pi_2 + 1$. To avoid biases, the two factors were unknown to us before factorization, as we generated the key via the RSA number generator in Ref. [36].

We emphasize that the RSA number N in Eq. (A27) is a representative instance of a 100-bit factorization problem, included for the purpose of reproducibility. Nonetheless, the TNSS algorithm successfully factorizes any 100-bit RSA key, and its effectiveness is not limited to this specific semiprime example reported.

The computational time was on the order of days on a single Intel(R) Xeon(R) Gold 5320 CPU node.

6. Analysis on the AsrPL scaling model

The data points in Fig. 2(a) are obtained by rescaling the average number of sr-pairs per lattice (AsrPL) ρ_{sr} by the number of samples per CVP, ℓ^γ , and the RSA input bit-length ℓ by an appropriate function of the number of qubits n . In this appendix, we show how we selected the rescaling functions and how we determined the optimal rescaling parameter $\omega = \omega^* \approx 8.3$.

The function AsrPL depends on both the number of qubits n and the bit-length ℓ . In the following, we analyze the average number of sr-pairs per lattice ρ_{sr} as a function of the number of qubits, rather than RSA bit-length. To this end, we test two alternative rescaling models

$$g_\omega(\ell) = \ell^\omega, \quad (\text{A29})$$

and

$$g_\xi(\ell) = \frac{1}{\Gamma_2} e^{\ell/\xi}, \quad (\text{A30})$$

with ω , Γ_2 , and ξ as fit parameters. As shown in Fig. 7, these two models can be clearly distinguished within the range of our simulation data.

To identify the more appropriate rescaling model, we perform a statistical analysis of the rescaled data. In Fig. 8(a), we plot the same simulation dataset for ρ_{sr} from Fig. 2(a), now




FIG. 7. Rescaling function comparison. Comparison between two functions $g_\omega(\ell) = \ell^\omega$ (blue curve) and $g_\xi(\ell) = \frac{1}{\Gamma_2} \exp(\ell / \xi)$ (orange curve), where $\omega = \omega^* = 8.3$, $\xi = \xi^* = 13$, and $\Gamma_2 \approx 275.52$, used to rescale the same data points of Fig. 2(a) but over the bit-length range. The optimal value of Γ_2 is computed through a least-squares procedure in $1 / \Gamma_2$ to minimize the relative difference between the two functions, while optimal rescaling parameter values are derived from the bootstrapped chi-squared curves. The left y axis is in log scale. The dashed gray curve shows the relative difference $\varepsilon_{\text{rel}}(\ell) = [g_\xi(\ell) - g_\omega(\ell)] / \{\frac{1}{2}[g_\omega(\ell) + g_\xi(\ell)]\}$ between the two functions.

as a function of the number of qubits n . Figure 8(b) shows the data after rescaling by $g_\omega(\ell)$ on the x axis and ℓ^γ on the y axis. The collapse onto a single curve persists across all instances explored, revealing a universal scaling behavior. To quantify the statistical uncertainty and evaluate which model better fits the data, we apply a *bootstrap* approach [54]. We fit the collapsed curves to a degree-2 polynomial for both rescaling schemes and perform 50 000 bootstrap resamplings of the rescaled data (excluding simulation errors) to estimate the distribution of chi-squared values. Note that, since each data point is assigned equal uncertainty, the resulting chi-squared values represent unweighted sums of squared residuals.

This procedure is repeated across a range of ω and ξ values near initial guesses to identify the minimum chi-squared value. The resulting Monte Carlo estimates of the average chi-squared and its statistical uncertainty are shown in Fig. 9(a). Our analysis demonstrates that the power-law rescaling consistently yields a lower (and thus better) chi-squared value than the exponential one.

To assess the confidence level in our model comparison, we compute an *empirical p value*. This quantifies how often the bootstrap procedure favors one rescaling over the other at their respective optimal parameters $\omega = \omega^* \approx 8.3$ and $\xi = \xi^* \approx 13$. Specifically, it provides a Monte Carlo estimate of probability $P(\chi^2[\omega^*] < \chi^2[\xi^*])$, that is, how often the power-law rescaling yields a lower chi-squared value than the exponential rescaling purely by chance [54]. The resulting p value indicates a strong preference for the power-law model, with a value of approximately 0.9949.




FIG. 8. Rescaling process. (a) The average number of sr-pairs per lattice (AsrPL) ρ_{sr} as a function of the number of qubits n . The different curves show the AsrPL for different values of the hyperparameter γ and RSA-key bit-length ℓ . Each data point represents an average over ten random RSA keys. Each ρ_{sr} is computed using $n_{\text{CVP}} = 50$ lattices. The y axis is plotted on a logarithmic scale. The red dashed line marks $\rho_{sr} = 1$, corresponding to one useful pair per CVP on average. (b) The same data as in panel (a), but rescaled: the ρ_{sr} values are rescaled by ℓ^γ on the y axis, and the number of qubits is rescaled by $g_{\omega^*}(\ell)$ on the x axis, with $\omega^* \approx 8.3$. Both the x axis and y axis are plotted on a logarithmic scale.

We further validate these findings using a second quantitative criterion to assess the quality of the data collapse under each rescaling model. Specifically, we employ the *Bhattacharjee-Seno* (BS) collapse quality score, a well-established metric for evaluating how well numerical simulation datasets collapse onto a single universal curve under a proposed rescaling hypothesis [55].

The BS score, denoted by S , measures the average distance between each curve and the local average of all others in its neighborhood. The width of the local neighborhood is controlled by a tunable parameter $\epsilon \in \mathbb{R}$, which defines the width of the local comparison window.

Using the previously identified optimal rescaling parameters from the bootstrapped chi-squared test, we now




FIG. 9. Statistical analysis and model comparison under rescaling transformations. (a) The bootstrapped chi-squared statistics [54] for evaluating the quality of the two rescaling models. For each model, the rescaled data are fitted using a degree-2 polynomial, and the corresponding chi-squared values are computed from 50 000 bootstrap samples across varying rescaling parameter values centered in an initial guess. The blue curve shows the result for the polynomial rescaling model, which exhibits a clear minimum at $\omega = \omega^* \approx 8.3$ with a value of $\chi_{\omega^*}^2 = 4.0 \pm 1.5$, indicating a well-defined optimal rescaling. In contrast, the orange curve for the exponential rescaling model displays a relatively flat profile with a minimum at $\xi = \xi^* \approx 13$ with $\chi_{\xi^*}^2 = 13.2 \pm 3.6$, suggesting a less robust optimal parameter. The empirical p value derived from the comparison strongly favors the polynomial rescaling model, with a value of approximately 0.9949. (b) The Bhattacharjee-Seno (BS) collapse quality score $S_R(\epsilon)$ [55] is plotted as a function of the neighborhood width parameter $\epsilon \in [0.05, 5.0]$. The blue curve corresponds to the polynomial rescaling model using the function $g_{\omega}(\ell)$, while the orange curve corresponds to the exponential rescaling model with $g_{\xi}(\ell)$. Both curves are computed using the optimal rescaling parameters $\omega^* = 8.3$ and $\xi^* = 13$, respectively, determined via (a) a bootstrap-based analysis of the chi-squared statistic. The y axis is plotted on a logarithmic scale.

compute the corresponding BS scores $S_{\omega^*}(\epsilon)$ and $S_{\xi^*}(\epsilon)$ over the range $\epsilon \in [0.05, 5]$. The results are shown in Fig. 9(b).

We find that the BS score is consistently lower for polynomial rescaling across the entire range of ϵ , indicating a more

coherent and tighter data collapse. This second independent statistical measure further supports our conclusion that the power-law rescaling $[g_{\omega}(\ell)]$ provides a more accurate and robust description of the scaling behavior than the exponential alternative $[g_{\xi}(\ell)]$.

- [1] R. L. Rivest, A. Shamir, and L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, *Commun. ACM* **21**, 120 (1978).
- [2] S. D. Galbraith, *Mathematics of Public Key Cryptography* (Cambridge University Press, Cambridge, England, 2012).
- [3] D. Aggarwal and U. Maurer, Breaking RSA generically is equivalent to factoring, *IEEE Trans. Inf. Theory* **62**, 6251 (2016).
- [4] R. Crandall and C. Pomerance, *Prime Numbers: A Computational Perspective*, 2nd ed. (Springer, New York, 2006).
- [5] K. Rabah, Review of methods for integer factorization applied to cryptography, *J. Appl. Sci.* **6**, 458 (2006).
- [6] F. Boudot, P. Gaudry, A. Guillevis, N. Heninger, E. Thomé, and P. Zimmermann, The state of the art in integer factoring and breaking public-key cryptography, *IEEE Security & Privacy* **20**, 80 (2022).
- [7] C. Pomerance, A Tale of Two Sieves, in *Biscuits of Number Theory*, edited by A. Benjamin and E. Brown (American Mathematical Society, Providence, RI, 2009), pp. 85–104.
- [8] P. Stevenhagen, The Number Field Sieve, in *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*, Mathematical Sciences Research Institute Publications, edited by J. P. Buhler and P. Stevenhagen (Cambridge University Press, Cambridge, England, 2008), Vol. 44, pp. 83–100.
- [9] T. Kleinjung, K. Aoki, J. Franke, A. K. Lenstra, E. Thomé, J. W. Bos, P. Gaudry, A. Kruppa, P. L. Montgomery, D. A. Osvik, H. te Riele, A. Timofeev, and P. Zimmermann, Factorization of a 768-Bit RSA modulus, in *Advances in Cryptology—CRYPTO 2010*, Lecture Notes in Computer Science, edited by T. Rabin (Springer, Berlin, 2010), Vol. 6223, pp. 333–350.
- [10] F. Boudot, P. Gaudry, A. Guillevis, N. Heninger, E. Thomé, and P. Zimmermann, Comparing the difficulty of factorization and discrete logarithm: A 240-digit experiment, in *Advances in Cryptology – CRYPTO 2020*, Lecture Notes in Computer Science, edited by D. Micciancio and T. Ristenpart (Springer International, Cham, Switzerland, 2020), Vol. 12171, pp. 62–91.
- [11] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Comput.* **26**, 1484 (1997).
- [12] C. Gidney and M. Ekerå, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, *Quantum* **5**, 433 (2021).
- [13] C.-P. Schnorr, Fast Factoring Integers by SVP Algorithms, corrected, Cryptology ePrint Archive, Report 2021/933 (2021).
- [14] L. Babai, On Lovász’ lattice reduction and the nearest lattice point problem, *Combinatorica* **6**, 1 (1986).
- [15] B. Yan, Z. Tan, S. Wei, H. Jiang, W. Wang, H. Wang, L. Luo, Q. Duan, Y. Liu, W. Shi, Y. Fei, X. Meng, Y. Han, Z. Shan, J. Chen, X. Zhu, C. Zhang, F. Jin, H. Li, C. Song, *et al.*, Factoring integers with sublinear resources on a superconducting quantum processor, [arXiv:2212.12372](https://arxiv.org/abs/2212.12372).
- [16] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero, The tensor networks anthology: Simulation techniques for many-body quantum lattice systems, *SciPost Phys. Lect. Notes* **8**, 8 (2019).
- [17] M. Ballarin, P. Silvi, S. Montangero, and D. Jaschke, Optimal sampling of tensor networks targeting wave function’s fast decaying tails, *Quantum* **9**, 1714 (2025).
- [18] C.-P. Schnorr, Factoring integers and computing discrete logarithms via diophantine approximation, in *Advances in Cryptology—EUROCRYPT ’91*, Lecture Notes in Computer Science Vol. 547, edited by D. W. Davies (Springer, Berlin, 1991), pp. 281–293.
- [19] P. van Emde-Boas, *Another NP-Complete Partition Problem and the Complexity of Computing Short Vectors in a Lattice*, Technical Report (Department of Mathematics, University of Amsterdam, Amsterdam, The Netherlands, 1981).
- [20] D. Micciancio, The hardness of the closest vector problem with preprocessing, *IEEE Trans. Inf. Theory* **47**, 1212 (2001).
- [21] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (N.Y.)* **326**, 96 (2011).
- [22] S. Montangero, *Introduction to Tensor Network Methods: Numerical Simulations of Low-Dimensional Many-Body Quantum Systems* (Springer, Cham, Switzerland, 2018).
- [23] S. V. Grebnev, M. A. Gavreev, E. O. Kiktenko, A. P. Guglya, A. R. Efimov, and A. K. Fedorov, Pitfalls of the sublinear QAOA-based factorization algorithm, *IEEE Access* **11**, 134760 (2023).
- [24] T. Khattar and N. Yosri, A comment on “Factoring integers with sublinear resources on a superconducting quantum processor”, [arXiv:2307.09651](https://arxiv.org/abs/2307.09651).
- [25] W. Aboumrar, D. Widdows, and A. Kaushik, Quantum and classical combinatorial optimizations applied to lattice-based factorization, [arXiv:2308.07804](https://arxiv.org/abs/2308.07804).
- [26] D. Coppersmith, Solving linear equations over GF(2): Block Lanczos algorithm, *Linear Algebra Appl.* **192**, 33 (1993).
- [27] S. Yang, An improvement of Babai’s rounding procedure for CVP, *Int. J. Netw. Secur.* **25**, 332 (2023).
- [28] A. K. Lenstra, J. Lenstra, H. W. Lenstra, Jr., and L. Lovász, Factoring polynomials with rational coefficients, *Math. Ann.* **261**, 515 (1982).
- [29] Y. Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).
- [30] Quantum TEA, Quantum tea—the quantum tensor network emulator applications (2025), Accessed: September 29, 2025, <https://www.quantumtea.it/>.
- [31] F. Baccari, D. Bacilieri, M. Ballarin, F. P. Barone, F. Campaioli, A. G. Catalano, G. Cataldi, A. Coppi, A. Costantini, A. Datta, A. D. Girolamo, D. Jaschke, S. B. Kožić, G. Magnifico, C. Mordini, S. Montangero, S. Notarnicola, A. Pagano, L. Pavesic,

- D. Rattacaso, *et al.*, Quantum tea: Qtealeaves (2025), <https://doi.org/10.5281/zenodo.14883066>.
- [32] G. V. Bard, *Algebraic Cryptanalysis* (Springer, New York, 2009).
- [33] C. Pomerance, Smooth numbers and the quadratic sieve, in *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*, MSRI Publications Vol. 44, edited by J. P. Buhler and P. Stevenhagen (Cambridge University Press, Cambridge, England, 2008), pp. 69–81.
- [34] L. Ducas, SchnorrGate, GitHub repository (2022), Accessed: June 14, 2024, <https://github.com/lducas/SchnorrGate>.
- [35] B. Kaliski, RSA factoring challenge, in *Encyclopedia of Cryptography and Security*, edited by H. C. A. van Tilborg (Springer, Boston, 2005), pp. 531–532.
- [36] N. Klein, Big Primes (2024), Accessed: June 14, 2024, www.bigprimes.org.
- [37] *The Development of the Number Field Sieve*, Lecture Notes in Mathematics, edited by A. K. Lenstra and H. W. Lenstra (Springer, Berlin, 1993), Vol. 1554.
- [38] J. Wagstaff, S. S., *The Joy of Factoring*, Student Mathematical Library (American Mathematical Society, Providence, RI, 2013), Vol. 68.
- [39] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [40] K. Blekos, D. Brand, A. Ceschini, C.-H. Chou, R.-H. Li, K. Pandya, and A. Summer, A review on quantum approximate optimization algorithm and its variants, *Phys. Rep.* **1068**, 1 (2024).
- [41] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 5 (2014).
- [42] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys. (N.Y.)* **349**, 117 (2014).
- [43] M. L. Wall and L. D. Carr, Out-of-equilibrium dynamics with matrix product states, *New J. Phys.* **14**, 125015 (2012).
- [44] P. C. G. Vlaar and P. Corboz, Simulation of three-dimensional quantum systems with projected entangled-pair states, *Phys. Rev. B* **103**, 205137 (2021).
- [45] J. Tindall, M. Fishman, E. Miles Stoudenmire, and D. Sels, Efficient tensor network simulation of IBM’s Eagle kicked Ising experiment, *PRX Quantum* **5**, 010308 (2024).
- [46] T. Felser, P. Silvi, M. Collura, and S. Montangero, Two-dimensional quantum-link lattice quantum electrodynamics at finite density, *Phys. Rev. X* **10**, 041040 (2020).
- [47] G. Magnifico, T. Felser, P. Silvi, and S. Montangero, Lattice quantum electrodynamics in (3+1)-dimensions at finite density with tensor networks, *Nat. Commun.* **12**, 3600 (2021).
- [48] S. Cavinato, T. Felser, M. Fusella, M. Paiusco, and S. Montangero, Optimizing radiotherapy plans for cancer treatment with tensor networks, *Physics in Medicine & Biology* **66**, 125015 (2021).
- [49] D. Jaschke, M. L. Wall, and L. D. Carr, Open source matrix product states: Opening ways to simulate entangled many-body quantum systems in one dimension, *Comput. Phys. Commun.* **225**, 59 (2018).
- [50] E. M. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, *New J. Phys.* **12**, 055026 (2010).
- [51] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Matrix product density operators: Simulation of finite-temperature and dissipative systems, *Phys. Rev. Lett.* **93**, 207204 (2004).
- [52] M. Zwolak and G. Vidal, Mixed-state dynamics in one-dimensional quantum lattice systems: A time-dependent super-operator renormalization algorithm, *Phys. Rev. Lett.* **93**, 207205 (2004).
- [53] A. H. Werner, D. Jaschke, P. Silvi, M. Kliesch, T. Calarco, J. Eisert, and S. Montangero, Positive tensor network approach for simulating open quantum many-body systems, *Phys. Rev. Lett.* **116**, 237201 (2016).
- [54] M. R. Chernick, *Bootstrap Methods: A Guide for Practitioners and Researchers*, 2nd ed., Wiley Series in Probability and Statistics (John Wiley & Sons, Hoboken, NJ, 2008).
- [55] S. M. Bhattacharjee and F. Seno, A measure of data collapse for scaling, *J. Phys. A: Math. Gen.* **34**, 6375 (2001).