

A Posteriori Error Analysis and Rank-Adaptive Structure-Preserving Numerical Integration of Large-Scale Lindblad Matrix Differential Equations

Authors Names	Abstract
Abdullah Issa Jassem^a Publication data: 21/9/ 2026 Keywords: Lindblad equation; a posteriori error; low-rank matrix differential equation; structure-preserving integration; rank adaptivity; trace norm.	Large-scale Lindblad matrix differential equations are difficult to integrate because a dense density matrix requires quadratic storage while standard time integrators can destroy positivity. This paper develops a second-order rank-adaptive structure-preserving integrating-factor method, denoted RA-SPIF2, together with a fully computable a posteriori error analysis in the trace norm. The update is derived from an explicit midpoint rule in the interaction picture and is implemented as a sum of positive semidefinite congruences acting directly on a tall factor. A thin QR factorization and a small Hermitian eigenproblem select the numerical rank. The discarded eigenvalue mass yields an exact trace-distance identity for the compression perturbation, while normalization enforces unit trace. For error control, two accepted half steps define a quadratic continuous reconstruction. Contractivity of the exact Lindblad semigroup gives a global residual bound, and a Bernstein representation converts each local residual integral into a certified expression involving only three trace norms. These norms are evaluated through small projected Hermitian matrices, without forming the density matrix or the Liouvillian. Numerical experiments for a driven thermal oscillator exhibit observed orders 1.74, 1.87, and 1.93 under refinement; the certified estimator remains reliable with effectivities from 22.36 to 17.49, while a Gauss residual estimator gives effectivities from 7.20 to 5.51. Rank studies reveal a clear temporal-error plateau, and coarse-step comparisons show positivity loss for classical RK4 but only roundoff-level violations for RA-SPIF2. A boundary-driven spin chain is advanced to dimension 16,384 with rank 10, using 2.5 MiB for the state factor instead of 4,096 MiB for a dense complex matrix.

1. Introduction

A typical example of Markovian open quantum dynamics is a (density) matrix $\rho(t)$ evolving through a (Lindblad) master equation. It is a formulation of coherent Hamiltonian motion in conjunction with irreversible action of the environment and it preserves the Hermiticity, positive semidefiniteness and unit trace. A dense matrix is almost always the main computational limitation of the current applications in quantum control [1]–[3] and noisy circuits [4]–[5] and also for reservoir engineering [6] and many-body simulations [7]–[10]. Hamiltonian simulation [16] and a learning approach based on simulation [17] further expand the computational spectrum of open-system models. The density matrix would have d squared complex elements and an explicitly constructed Liouvillian would be even bigger if the Hilbert-space dimension is d .

The density matrix has a physical constraint, it isn't cosmetic. Negative eigenvalues are incompatible with the probabilistic character of the state and trace drift causes the expectation values to be incorrect and may cause long simulations to become unstable. Deterministic schemes that preserve the structure of the system are, therefore, in the spotlight. Cao and Lu developed arbitrary order schemes avoiding nonphysical states [4] and Appelö and Cheng have developed high order low rank schemes based on Kraus representations [5]. These improvements demonstrate that for both accuracy and physical admissibility, a "one size fits all" discretization does not work, but has to be designed with each other in mind.

A second path to scalability are low-rank state representations. If the state of the system is in moderate entropy, it is possible to write $\rho_n = Y_n Y_n^\dagger$ with $Y_n \in \mathbb{C}^{d \times r_n}$ and $r \ll d$. Adaptive variational evolution [6], ensemble truncation [7] [8] and robust dynamical low-rank integrators [9]–[13] have shown that the rank should evolve with the physical mixedness and not be specified, a priori. But a rank judgement using singular values alone does not take into account the total time-discretisation error, and a standard embedded estimate does not necessarily uphold the density-matrix structure.

There is an added challenge of certification. If that estimate is going to be needed - the only reason for providing an error for a large matrix differential equation is if the error estimate can be calculated using the same low rank data that the integrator was used with. The advantage of the residual-based a posteriori analysis is that it can remove the same equation from the numerical trajectory and thus have varying steps [18]. It appears it is the only way to be able to make a point-wise diagonalisation of a dense Hermitian matrix to be able to make a direct residual evaluation in the trace norm. This would make the factor-representation benefit moot. The principle remark of the present work is that all the "residual" terms came by the Lindblad operator and the quadratic reconstruction have a compact indefinite factorisation. A projected matrix whose size is proportional to the numerical rank and number of jump operators can therefore be used to recover its nonzero eigenvalues from.

In this paper the ideas get consolidated into a single structure. The main findings are the following:

- (1) An explicit sum of positive semidefinite congruences is written in the form of a second order integrating factor midpoint step. The factors are implemented in batch with batched exp actions and never a full density matrix is formed.
- (2) An exact formula for the trace distance is available for a rank-adaptive QR-eigenvalue compression. The formula converts wasted spectral mass to a physical measurement of perturbation loss that is transparent and distinguishes compression from trace drift.
- (3) The same two accepted half steps can be used to get a global trace-norm a posteriori estimate, reconstructed using a quadratic interpolation. With a Bernstein-polynomial argument, we have a certified local upper bound, and – only – 3 residual norms.
- (4) A small hermitian eigenproblem is used to assess each of the residual trace norms. This allows using the estimator for sparse operators, curly exponential actions and dynamically changing rank.

Convergence and other results (estimator-effectivity, rank-error-balancing, positivity and memory scalings) are numerically tested on oscillator and spin-chain models.

The created method is referred to as the RA-SPIF2 (rank-adaptive structure-preserving an integrating-factor method of order two). The compressed state update is nonlinear indeed, due to the selection of the rank and the normalization and therefore it is called here state-structure preserving instead of a linear completely positive trace-preserving channel. The raw update is of course a totally positive linear map and the accepted numerical state is in the density-matrix set by construction.

The rest of the formulas are arranged as below. The matrix model is introduced in section 2, and the low-rank case is considered as a special case. In Section 3, RA-SPIF2 is derived and its consistency and structure properties are proven. In Section 4 the a posteriori estimates of the trace-norms and their reduced evaluation are realized. The adaptive algorithm is provided in section 5. It has also numerical experiments in Section 6 and limitation as well as conclusion in Section 7 and Section 8 respectively.

2. Lindblad Matrix Dynamics and Low-Rank Setting

2.1. Governing equation and physical state space

Let $H = H^\dagger$ be the Hamiltonian and let L_ℓ denote the jump operators. The autonomous Lindblad equation is

$$\begin{aligned}\dot{\rho}(t) &= -i[H, \rho(t)] \\ &+ \sum_{\ell=1}^m \left(L_{\ell} \rho(t) L_{\ell}^{\dagger} - \frac{1}{2} \{ L_{\ell}^{\dagger} L_{\ell}, \rho(t) \} \right), \\ \rho(0) &= \rho_0.\end{aligned}\tag{1}$$

Here the commutator is responsible for coherent motion, anticommutator for damping and jump terms for environmental transitions. With the introduction of the effective non-Hermitian generator and the jump map, the equivalent split form is introduced.

$$\begin{aligned}A &= -iH - \frac{1}{2} \sum_{\ell=1}^m L_{\ell}^{\dagger} L_{\ell}, \\ \mathcal{J}(X) &= \sum_{\ell=1}^m L_{\ell} X L_{\ell}^{\dagger}, \\ \mathcal{L}(X) &= AX + XA^{\dagger} + \mathcal{J}(X).\end{aligned}\tag{2}$$

The physically admissible states form the convex set $\mathcal{D}_d$, defined by

$$\mathcal{D}_d = \{X \in \mathbb{C}^{d \times d} : X = X^{\dagger}, X \geq 0, \text{tr} X = 1\}.\tag{3}$$

For a time-independent Lindbladian, the exact propagator is a completely positive trace-preserving semigroup. A property used throughout the analysis is trace-norm contractivity on Hermitian perturbations:

$$\begin{aligned}\|\Phi_t(X) - \Phi_t(Y)\|_1 &\leq \|X - Y\|_1, \\ X &= X^{\dagger}, \quad Y = Y^{\dagger}, \quad t \geq 0.\end{aligned}\tag{4}$$

One direct motivation can be viewed in the light of the Jordan decomposition of a Hermitian matrix in terms of a positive and negative part. The first map, to positivity, takes those parts into positive matrices; the second rule of preservation is that this preserves the traces; and the third rule, the triangle, gives the contraction. This argument avoids assumptions about diagonalizability of the superoperator.

2.2. Low-rank representation and matrix-free assumptions

The numerical states at time t_n are represented by the numerical state density operator of the form $\rho_n = Y_n Y_n^{\dagger}$. The columns of the factor do not need to be orthogonal, but it is expected that the rank $r < d$ of the factor is moderate for rapidly decaying spectrum of the state. The algorithm requires the presence of products with the same $H = H^{\dagger}$ and L_{ℓ} and these are most readily provided by either sparse matrices or matrix-free operator routines. The exponential actions are computed on a block of vectors using a Krylov or polynomial method without ever constructing the exponential matrix $E(s) = \exp(sA)$.

The difference between state rank and the complexity of the operator powers them appreciably. The Hamiltonian density is not necessarily low cost if it is low rank. The one that is intended is thus a structured or sparse effective generator with moderate state entropy. The sparse Kronecker operators are utilized in the spin-chain tests of Section 6, while banded ladder operators are utilized in the oscillator tests.

2.3. Position with respect to recent methods

Various aspects of the computational problem are discussed in recent literature. Structure-preserving dense schemes pay attention to the physical admissibility [4]; Kraus-based low-rank integration yields high-order completely positive methods [5]; adaptive variational methods do adapt the representation space [6], [15], [22]; and generic dynamical low-rank (gLRA) methods offer robust rank adaptation

for matrix differential equations [9] to [14]. The current measurement is deliberately smaller so that there is a residual certificate of the physically relevant trace norm.

Table 1. Comparison of representative approaches and the present framework.

Approach	Low-rank	Adaptive rank	Density-matrix structure	A posteriori control
Cao–Lu [4]	Optional	No	Yes	No residual certificate
Appelö–Cheng [5]	Yes	Truncation based	CPTP-oriented	No trace-norm certificate
Gravina–Savona [6]	Yes	Yes	Physical ansatz	Variational diagnostics
Generic adaptive DLR [10]–[14]	Yes	Yes	Problem dependent	Local low-rank/time estimators
RA-SPIF2 (this work)	Yes	Yes	PSD and unit trace	Certified global trace-norm residual bound

The alleged novelty is the fact that the combination of the three features (i) a congruence-form low rank step, (ii) the exact compression-distance identity and (iii) a globally valid residual estimate whose trace norms are small matrices is claimed to be novel. A second order construction can be used here and results in the following particularly transparent reconstruction and estimator, although higher order extensions can be made by replacing the midpoint stage with a nonnegative coefficient integrating factor Runge–Kutta construction.

3. Rank-Adaptive Structure-Preserving Integrator

3.1. Interaction-picture midpoint construction

Set $E(s) = \exp(sA)$. Over one step of length h , define the interaction-picture state

$$X_n(s) = E(-s)\rho(t_n + s)E(-s)^\dagger, \quad 0 \leq s \leq h. \quad (5)$$

Differentiation cancels the stiff left and right actions generated by A and gives

$$\begin{aligned} \dot{X}_n(s) &= \mathcal{J}_s(X_n(s)), \\ \mathcal{J}_s(X) &= \sum_{\ell=1}^m B_\ell(s)XB_\ell(s)^\dagger, \\ B_\ell(s) &= E(-s)L_\ell E(s). \end{aligned} \quad (6)$$

The right-hand side is a sum of positive congruences at every interaction time. Applying the explicit midpoint rule to this nonautonomous equation gives

$$\begin{aligned} X_{n,\star} &= X_n(0) + \frac{h}{2}\mathcal{J}_0(X_n(0)), \\ X_{n+1} &= X_n(0) + h\mathcal{J}_{h/2}(X_{n,\star}). \end{aligned} \quad (7)$$

Transforming the second stage back to the physical picture yields the raw update

$$\begin{aligned}
\rho_{n+1}^{\text{raw}} &= E(h)\rho_n E(h)^\dagger + h \sum_{\ell=1}^m C_\ell(h) \Xi_n C_\ell(h)^\dagger, \\
\Xi_n &= \rho_n + \frac{h}{2} \sum_{j=1}^m L_j \rho_n L_j^\dagger, \\
C_\ell(h) &= E(h/2) L_\ell E(h/2).
\end{aligned} \tag{8}$$

All the terms of (8) are positive semidefinite if the input is positive semidefinite. The formula is not just an algebraic rewriting: it precisely uncovers the structure required for a factor implementation, and removes subtraction of positive density-matrix parts, from negative density-matrix parts.

3.2. Tall-factor implementation

Suppose $\rho_n = Y_n Y_n^\dagger$. Factor the midpoint argument and transport it to the physical midpoint through

$$\begin{aligned}
S_n &= \left[Y_n, \sqrt{\frac{h}{2}} L_1 Y_n, \dots, \sqrt{\frac{h}{2}} L_m Y_n \right], \\
W_n &= E(h/2) S_n.
\end{aligned} \tag{9}$$

The raw endpoint can then be assembled as one tall block

$$\begin{aligned}
Z_{n+1} &= [Z_n^{(0)}, Z_n^{(1)}, \dots, Z_n^{(m)}], \\
Z_n^{(0)} &= E(h) Y_n, \\
Z_n^{(\ell)} &= \sqrt{h} E(h/2) L_\ell W_n, \quad 1 \leq \ell \leq m, \\
\rho_{n+1}^{\text{raw}} &= Z_{n+1} Z_{n+1}^\dagger.
\end{aligned} \tag{10}$$

All jump operator products could be glued together before second half step exponential action. So it takes 3 block exponential actions: one for midpoint factor, one for homogenic end point and all carried across jump blocks. The dense state matrix is not created.

3.3. Rank selection and exact compression distance

The number of columns in the raw factor grows after each step and must be compressed. A thin QR factorization followed by an eigenvalue decomposition of the small Gram matrix gives

$$\begin{aligned}
Z_{n+1} &= Q_n R_n, \\
R_n R_n^\dagger &= V_n \text{diag}(\lambda_1, \dots, \lambda_k) V_n^\dagger, \\
\lambda_1 &\geq \dots \geq \lambda_k \geq 0.
\end{aligned} \tag{11}$$

Let τ_r be a relative discarded-mass tolerance. The accepted rank is the smallest one satisfying

$$\begin{aligned}
s &= \sum_{j=1}^k \lambda_j, \quad \delta_r = \sum_{j=r+1}^k \lambda_j, \\
r_{n+1} &= \min \left\{ r: \frac{\delta_r}{s} \leq \tau_r \right\}.
\end{aligned} \tag{12}$$

The normalized accepted factor is

$$\begin{aligned}
Y_{n+1} &= \frac{Q_n V_n(:, 1:r_{n+1}) \text{diag}(\sqrt{\lambda_1}, \dots, \sqrt{\lambda_{r_{n+1}}})}{\sqrt{s - \delta_{r_{n+1}}}}, \\
\rho_{n+1} &= Y_{n+1} Y_{n+1}^\dagger.
\end{aligned} \tag{13}$$

Proposition 1 (exact compression perturbation).

Ensure that the trace of the tall factor is positive, and keep the leading components in the spectral truncation of said factor as per QR–eigenvalue compression. Then the trace norm of the normalized raw state and the normalized truncated state is exactly

$$\left\| \frac{ZZ^\dagger}{s} - \frac{Y_r Y_r^\dagger}{s - \delta_r} \right\|_1 = \frac{2\delta_r}{s} \equiv \chi_r. \quad (14)$$

Proof.

Both the two normalized matrices have the same orthonormal eigenbasis. On the ‘retained subspace’ the addition of renormalization is an absolute shift in mass equal to the ‘discarded subspace’, the shift from the spectral tail is of the same size on the ‘discarded subspace’. They sum up to the right hand side of (14).

The exact identity is better than the Frobenius-tail heuristic because it quantifies the change in the operational trace norm. Also, it explains the reason of the tabulated compression indicator, which is around twice the relative discarded-mass tolerance. The transformation of the unnormalized raw matrix to the accepted one is computed by ensuring that it satisfies

$$\left\| ZZ^\dagger - \frac{Y_r Y_r^\dagger}{s - \delta_r} \right\|_1 \leq |s - 1| + \frac{2\delta_r}{s}. \quad (15)$$

The two components in (15) can be traced to different sources: the first term is the local trace defect of the raw time integrator, the second term represents rank truncation after normalization. They’re both cheap diagnostic tests. Therefore, they are not added on to the certified global bound; the residual estimator added later already takes them into account in the accepted trajectory.

3.4. Structure and order

Theorem 1 (structure preservation and second-order consistency). Let the generator and jump operators be bounded operators, and assume that for the exponential actions they are either exact or obtained to third order accuracy or better in the step size, in the local error. The map in (8) is positive semidefinite preserving, linear, completely positive and Hermiticity preserving for every input density matrix. The compression and normalization of the QR–eigenvalue in (13) results in a state that is in the density-matrix set. The normalized method is second order accurate without compression. Rank compression preserves 2nd order convergence on a fixed interval as long as the accumulated compression perturbation is not larger in order than the accumulated local time error; 3rd order local trace-distance perturbation of the step size is enough.

Proof. The raw map is positive, Hermiticity preserving and is a sum of positive congruences. From this follows that the factorization (13) is explicitly a normalized PD one. In the interaction picture, (7) is the explicit third order defect midpoint rule for the nonautonomous right hand side which is assumed to be continuously differentiable. The exact solution is at the same time trace preserving and, therefore, the defect of the raw trace and caused by normalization on it have the same order. The stability contribution comes from the contractility of the exact Lindblad semigroup and the compression contribution is given by Proposition 1.

The theorem makes an explicit distinction between the initial linear map and the resulting map of the appropriate rank. The latter is nonlinear and is dependent on the current spectrum. The point of it is to keep the state constraints -- not to construct a linear channel that acts on some operators.

3.5. Computational complexity

For m jump operators and rank r the raw factor width is $k = (m^2 + m + 1)r$. The leading algebraic costs are summarized by

$$\begin{aligned} k &= (m^2 + m + 1)r, & \mathcal{C}_{\text{mem}} &= \mathcal{O}(dk), \\ \mathcal{C}_{\text{QR}} &= \mathcal{O}(dk^2), & \mathcal{C}_{\text{eig}} &= \mathcal{O}(k^3). \end{aligned} \quad (16)$$

When the operator product is sparse the cost will be proportional to the nonzero count of the Hamiltonian and the jump operators times width of the block. Exponential actions prevail when the effective generator has high spectral radius or is stiff. They are linear with regard to the Hilbert-space dimension in each sparse matrix–block multiplication but vary in cost according to the Krylov dimension and tolerance chosen. In comparison, the dense state storage scales squared in this dimension, while an explicit construction of the Liouvillian is in general impractical.




Figure 1. Data flow of one residual-certified rank-adaptive macro step. All large objects are tall factors or sparse operator blocks; eigendecompositions are restricted to projected matrices.

4. A POSTERIORI TRACE-NORM ERROR ANALYSIS

4.1. Continuous quadratic reconstruction

Image analysis methods that are a posteriori are not based on nodal times, but rather a continuous path. Two accepted RA-SPIF2 half steps will produce one macro interval. Physical midpoint is from the 1st endpoint and macro endpoint is from the 2nd endpoint. For the accepted states normalized, give the quadratic reconstruction.

$$\begin{aligned} P_n(s) &= \ell_0(s)\rho_n + \ell_m(s)\rho_{n+1/2} + \ell_1(s)\rho_{n+1}, \\ s &= \frac{t - t_n}{h_n}. \end{aligned} \quad (17)$$

$$\ell_0(s) = 2s^2 - 3s + 1, \quad \ell_m(s) = -4s^2 + 4s, \quad \ell_1(s) = 2s^2 - s. \quad (18)$$

All 3 accepted states are exactly interpolated in the reconstruction. It expands to be Hermitian since its coefficients are real, but not necessarily positive between nodes. The above-mentioned reconstruction positivity is not a requirement for the residual estimate; only Hermiticity is required for ensuring the trace-norm contractivity of the residual equation.

The local residue of $P_n(t)$ is

$$\mathcal{R}_n(t) = \dot{P}_n(t) - \mathcal{L}(P_n(t)). \quad (19)$$

4.2. Global residual bound

Theorem 2 (global a posteriori estimate).

Use the same Hermitian initial state for the exact and the piecewise quadratic reconstruction. Then, at each final node the trace-norm error is

$$\begin{aligned} \|\rho(T) - P(T)\|_1 &\leq \|\rho_0 - P(0)\|_1 \\ &+ \sum_{n=0}^{N-1} \int_{t_n}^{t_{n+1}} \|\mathcal{R}_n(t)\|_1 dt. \end{aligned} \quad (20)$$

Proof. For every interval, the difference between the reconstruction equation and the exact Lindblad equation provides the evolution of the error on the interval. The local estimate is obtained from variation of constants, trace-norm contractivity of the exact propagator on Hermitian perturbations and the triangle inequality. By adding them all up over all the intervals, one can see that (20) can be proved.

The estimate is an unconditional estimate in regard to the production of the nodal states. For instance, rank changes, rejected trials, normalization and Krylov tolerances are represented in the residual of the accepted reconstruction. A separated proofs of consistency are not required for the validity of the a posteriori bound.

Bounded observables are also controlled with the trace norm. According to Hölder duality, a duality between the linear spaces associated with different observables O , yields for each Hermitian observable O ,

$$\left| \text{tr} \left(O(\rho(T) - P(T)) \right) \right| \leq \|O\|_\infty \|\rho(T) - P(T)\|_1. \quad (21)$$

Hence, all bound observables with known operator norm by the same residual certificate. This is why in the density-matrix calculations one prefers the trace norm over the Frobenius norm.

4.3. Bernstein residual certificate

While it is possible to directly numerically quadrature $\|\mathcal{R}_n(t)\|_1$, a guaranteed local upper bound follows from the Bernstein form. Following the pattern of poebasedir/dofonuclear.cpp, define the three Hermitian matrices:

$$\begin{aligned} B_0 &= \frac{-3\rho_n + 4\rho_{n+1/2} - \rho_{n+1}}{h_n} - \mathcal{L}(\rho_n), \\ B_1 &= \frac{-\rho_n + \rho_{n+1}}{h_n} + \frac{1}{2}\mathcal{L}(\rho_n) - 2\mathcal{L}(\rho_{n+1/2}) + \frac{1}{2}\mathcal{L}(\rho_{n+1}), \\ B_2 &= \frac{\rho_n - 4\rho_{n+1/2} + 3\rho_{n+1}}{h_n} - \mathcal{L}(\rho_{n+1}). \end{aligned} \quad (22)$$

With these coefficients, the residual polynomial is

$$\mathcal{R}_n(s) = B_0(1-s)^2 + 2B_1s(1-s) + B_2s^2. \quad (23)$$

The Bernstein basis is nonnegative on the unit interval and each basis function integrates to one third. Consequently,

$$\begin{aligned} \int_{t_n}^{t_{n+1}} \|\mathcal{R}_n(t)\|_1 dt &\leq \eta_n^B, \\ \eta_n^B &:= \frac{h_n}{3} (\|B_0\|_1 + \|B_1\|_1 + \|B_2\|_1). \end{aligned} \quad (24)$$

Thereby, $\sum_{n \in \mathcal{I}_{\text{acc}}} \eta_n^B$ over the accepted intervals is a certified upper bound in Theorem 2. Here is a less conservative practical indicator with n_q Gauss–Legendre points:

$$\eta_n^Q := \frac{h_n}{2} \sum_{q=1}^{n_q} w_q \left\| \mathcal{R}_n \left(\frac{1 + \xi_q}{2} \right) \right\|_1. \quad (25)$$

Quadratures remainders if they are not bounded are not certified and thus, the quadrature indicator is not certified. However, it is convenient for the step selection due to its smoothness and low dimensional integrand after projection. In the numerical section, the certified Bernstein estimator, and the quadrature estimator of the practical effectivity, are reported.

4.4. Reduced trace-norm evaluation

For a factor $Y_n \in \mathbb{C}^{d \times r_n}$, the Lindblad image has an indefinite low-rank representation

$$\begin{aligned} \mathcal{L}(YY^\dagger) &= Z_{\mathcal{L}} M_{\mathcal{L}} Z_{\mathcal{L}}^\dagger, \\ Z_{\mathcal{L}} &= [Y, AY, L_1 Y, \dots, L_m Y]. \end{aligned} \quad (26)$$

The coefficient matrix is such that the diagonal blocks of the jump factors are identity and the off diagonal blocks of the state factor with its image under the effective generator are identity. Any linear combination of nodal states and lindblad images is constructed by mimicking the stride of the factor blocks to concatenate them and storing the scalar coefficients in a block diagonal Hermitian matrix.

Proposition 2 (small-matrix trace norm). Suppose a Hermitian residual is factorized using above low rank factorization method and its tall factor is a thin QR factorized matrix. Then the exact value of the residual trace norm is calculated, up to floating-point round-off error, from the projected matrix:

$$\begin{aligned} \|ZMZ^\dagger\|_1 &= \|R_Z M R_Z^\dagger\|_1, \\ \|R_Z M R_Z^\dagger\|_1 &= \sum_j |\lambda_j(R_Z M R_Z^\dagger)|, \\ Z &= Q_Z R_Z, \quad M = M^\dagger. \end{aligned} \quad (27)$$

Proof. The nonzero eigenvalues of the full residual and the projected Hermitian matrix are the same, since the columns of the orthogonal factor are orthonormal, the full residual is zero on the orthogonal complement. The fact that the trace norm of a Hermitian matrix equals sum of the absolute values of the eigenvalues can be used to see why the claim is true.

This reduction is the “scalability” conclusion for the estimator. It is an alternative that substitutes a full-state Hermitian eigensolve by an eigensolve in a much smaller, projected state-space of dimension proportional to the concatenated rank. All three of the Bernstein coefficients and each of the quadrature residuals are evaluated in the same routine. The thin QR factorization automatically removes numerically dependent factor columns.

5. Residual-Certified Rank and Step Adaptivity

5.1. Error-budget allocation

Let them be denoted by TOL, and let the tolerance density be defined as $\varepsilon = \text{TOL}/T$. A conservative acceptance rule is that it will be no more than an allowable amount, $\eta = \varepsilon$ times the current step length, where $\varepsilon = \text{TOL}/T$ is a small fixed number. Selecting a fraction θ_r from the same local budget is allocated for the rank perturbation.

$$\begin{aligned}
\tau_{r,n} &= \frac{\theta_r \varepsilon h_n}{4}, & \varepsilon &= \frac{\text{TOL}}{T}, \\
\chi_{n,j} &\leq \frac{\theta_r \varepsilon h_n}{2}, & j &= 1, 2, \\
\chi_{n,1} + \chi_{n,2} &\leq \theta_r \varepsilon h_n.
\end{aligned} \tag{28}$$

There are two accepted compressions in a macro step. Proposition 1 thus puts an upper bound on each compression distance and on their sum as the full rank fraction displayed in (28). This allocation is not an extra term in the stringent estimator, but is rather a safeguard provided by the controller as a result of the rank truncation in the reconstruction residual already has the total effect of rank truncation. The rank tolerance is also bounded between upper and lower limits set by the user, to avoid sudden jumps in the case where the residual becomes very small.

5.2. Step controller

The global error of a second order method is $\mathcal{O}(h^3)$ and the local error density is $\mathcal{O}(h)$. The proper proportional controller then is the following:

$$h_{n+1} = \text{clip} \left(h_{\min}, h_{\max}, \theta_h h_n \left(\frac{\varepsilon h_n}{\eta_n} \right)^{1/2} \right). \tag{29}$$

clip optionally provides lower and upper bounds on steps and ratio, as well as a safety factor, θ_h . However, the certified mode is $\eta = \eta_n^B$, whereas the faster practical mode is $\eta = \eta_n^Q$ but with still reporting an accumulated certified estimator. Step rejection means that the factor remains unchanged.

Table 2. Rank-adaptive, residual-certified RA-SPIF2 macro step.

Step	Operation
1	Form the stage factor from the current state factor and apply the half-step exponential action.
2	Construct the raw endpoint factor by two batched exponential actions; do not form the density matrix or Liouvillian.
3	Compute a thin QR factorization of the raw factor and diagonalize the resulting small Hermitian matrix.
4	Choose the smallest rank satisfying the discarded-mass budget, normalize the retained factor, and record the exact compression distance and raw trace defect.
5	Repeat Steps 1–4 for a second half step; retain the accepted midpoint and endpoint factors.
6	Build the quadratic reconstruction, evaluate the three Bernstein residual coefficients in low-rank form, and compute the certified estimator; optionally compute the quadrature estimator.
7	Accept when the chosen local residual indicator lies within its allocated budget; otherwise reject the trial, reduce the step, and repeat.
8	Update the step size with the residual controller and update the next rank tolerance from the allocated rank-error fraction.

One factor-compression routine and one residual-norm routine can be used to implement the algorithm. The same QR kernel is used in both. Exponential action tolerances should therefore be less than the allocated local residual budget, otherwise we will see Krylov errors in the reconstruction

residual and this will lead to unnecessary step reduction. A series of fixed-step studies online seek to isolate the temporal convergence and rank effects from the same study, with a separate tolerance study to engage the residual controller.

6. Numerical Experiments

6.1. Implementation and diagnostics

The prototype was realized in with Python (in this case version 3.13.5) and the NumPy libraries (version 2.3.5), SciPy (version 1.17.0). Sparse matrix–block products and exponential actions were carried out without compiling the Liouvillian. All runs were reported with five virtual Intel Xeon Platinum 8573C CPU cores and 5.9 GiB of memory. Wall-clock is single-run on a common environment, and should instead be interpreted with the algorithmic dimensions and not as a hardware-independent measurement.

For test sizes smaller than 128, the reference state was generated by taking the exponential of the vectorized Lindblad operator and acting it on the initial density matrix. The main error diagnosis and effectivity tests were

$$\begin{aligned} e_1(T) &= \|\rho_{\text{num}}(T) - \rho_{\text{ref}}(T)\|_1, \\ \mathcal{E}(\eta) &= \frac{\eta}{e_1(T)}, \quad \mathcal{P} = \text{tr}(\rho^2). \end{aligned} \quad (30)$$

Test Physical structure with smallest eigen value and smallest trace defect value. Only a numerical state was saved, in the form of a factor state; a dense state was created just for diagnostics and reference purposes for small-problems.

6.2. Driven thermal oscillator: convergence and estimator effectivity

The first test is a squeezed (truncated) driven harmonic oscillator with annihilation operator, frequency of the oscillator, drive amplitude, damping rate and thermal occupation:

$$\begin{aligned} H &= \omega a^\dagger a + \Omega(a + a^\dagger), \\ L_1 &= \sqrt{\kappa(n_{\text{th}} + 1)} a, \quad L_2 = \sqrt{\kappa n_{\text{th}}} a^\dagger. \end{aligned} \quad (31)$$

Note that the initial state is the 3rd Fock state. A 24-level truncation is employed, and the intrinsic frequency of the oscillators as well as the amplitude and the damping rate of the oscillators are 1, 0.6, and 0.35 respectively. The temperature and the final time in the convergence test are 0.25 and 2. The discarded mass tolerance is 10^{-13} , and the temporal error dominates. A half step of the RA-SPIF2 and contains two macro steps which supply the reconstruction point in the middle of the two.

Table 3. Trace-norm convergence and residual estimator effectivity for the driven oscillator.

Step size	Error	Order	Quadrature	Quad. eff.	Certified	Cert. eff.	Max rank
0.400	3.2981×10^{-2}	—	2.3740×10^{-1}	7.20	7.3743×10^{-1}	22.36	18
0.200	9.8617×10^{-3}	1.74	6.1309×10^{-2}	6.22	1.9451×10^{-1}	19.72	20
0.100	2.7048×10^{-3}	1.87	1.5527×10^{-2}	5.74	4.9350×10^{-2}	18.25	21
0.050	7.0842×10^{-4}	1.93	3.9064×10^{-3}	5.51	1.2390×10^{-2}	17.49	21




Figure 2. Final trace-norm error, practical residual estimator, certified Bernstein estimator, and a second-order reference slope.

As the step is refined, the observed order goes towards two. Both the estimators go down in a similar rate of decline as the error. The Bernstein certificate is made in this way to be obviously conservative, and changes the norm of a weighted sum into the weighted sum of norms, its effectivity is decreased from 22.36 to 17.49. Quadrature residual is significantly more pointed with effectivity dropping from 7.20 to 5.51. Both indicators are not focused on the problem, and both are well above the measured final error.

6.3. Certified adaptive-step validation

The residual controller was next activated on the same oscillator. A safety factor of 0.85 and initial macro step 0.4 were used, while the rank tolerance was fixed at 10^{-12} to isolate time adaptation. Every accepted step satisfied the certified local residual budget, so the accumulated estimator was bounded by the requested global tolerance.

Table 4. Certified adaptive-step performance for the driven oscillator.

Tolerance	Accepted	Rejected	Step-size range	Final error	Certified estimator	Effectivity
0.200	12	2	0.0799–0.3322	1.3705×10^{-2}	9.7141×10^{-2}	7.09
0.100	16	2	0.0555–0.2504	9.0379×10^{-3}	5.5099×10^{-2}	6.10
0.050	22	2	0.0387–0.1915	5.1805×10^{-3}	2.9592×10^{-2}	5.71
0.020	34	2	0.0242–0.1294	2.3312×10^{-3}	1.2619×10^{-2}	5.41
0.010	47	3	0.0170–0.0943	1.3240×10^{-3}	6.6703×10^{-3}	5.04

The cumulative certificate and the actual final error are not exceeded by all the requested tolerances. The tighter the tolerance, the more steps can be accepted, in this case 12 to 47, and the controller will only need 2 of 3 steps to be rejected. The reduction in effectivity is due to a greater localisation on the adaptive intervals which are smaller, and the certificate exhibits the same good reliability without specific tuning to the problems at hand.

6.4. Rank adaptivity and the temporal-error plateau

In order to isolate the rank behavior, the number of levels to truncate the given oscillator is raised to 32 levels, and it is integrated to the final time 4 with macro step 0.1. The relative discarded-mass tolerance is varied across Table 5. The maximum compression-distance column provides the actual perturbation (according to Proposition 1) when the perturbation is maximum along steps that are accepted.

Table 5. Effect of the relative discarded-mass tolerance on rank and final error.

Rank tolerance	Final rank	Max rank	Max compression distance	Final error
10^{-3}	6	6	1.753×10^{-3}	1.5869×10^{-2}
10^{-4}	7	7	1.967×10^{-4}	6.7964×10^{-3}
10^{-5}	9	9	1.979×10^{-5}	4.6164×10^{-3}
10^{-6}	10	10	1.972×10^{-6}	4.5040×10^{-3}
10^{-8}	13	13	1.904×10^{-8}	4.4848×10^{-3}
10^{-10}	16	16	1.997×10^{-10}	4.4847×10^{-3}
10^{-12}	19	19	1.997×10^{-12}	4.4847×10^{-3}




Figure 3. Final error versus rank tolerance. Below approximately 10^{-8} , the time-discretization error dominates and further rank growth gives no visible accuracy benefit.

The results show a balance that is the same as one would expect. Loose tolerances can't be seen and they degrade the rank of the solution. When the rank tolerance is at 10^{-8} , the error is found to remain almost at 4.4847×10^{-3} and the final rank continues to rise from 13 to 19. A compression error must therefore be kept fairly high in an adaptive code in order not to drive the error below the minimum currently attained in the temporal residual. The precise compression-distance indicator quantitative-that balancing decision.

6.5. Positivity under a coarse step

For the deliberately coarse comparison, a 16-level oscillator, drive amplitude 0.4, damping rate 0.3, thermal occupation 0.15, final time 2 and six macro steps are used. Classical dense RK4 is not positivity preserving for this linear problem; however, it preserves trace. The rank tolerance of RA-SPIF2 is 10^{-13} .

Table 6. Coarse-step structure and accuracy comparison.

Method	Final trace-norm error	Minimum eigenvalue	Maximum trace defect
RA-SPIF2	5.9329×10^{-3}	-2.56×10^{-22}	8.88×10^{-16}
Classical RK4	1.4604×10^{-1}	-2.0877×10^{-2}	0.00




Figure 4. Smallest eigenvalue at the accepted macro nodes. RA-SPIF2 remains nonnegative up to roundoff, whereas classical RK4 develops a substantial negative eigenvalue.

There are two structural properties being compared. The reason for not obtaining zero trace from RK4 is because the trace of the semidiscrete right-hand side is also zero, but RK4 returns a negative eigenvalue of approximately 2.1×10^{-2} and a definitely larger final error. RA-SPIF2 performs a unit trace enforcing the positive semidefiniteness to floating-point accuracy. The value in Table 6 which is very small but negative is due to the round off error of the eigensolver but it does not reflect the loss of positivity.

6.6. Boundary-driven spin-chain scaling

The large-scale test is an anisotropic nearest-neighbor chain of q spins, with $d = 2^q$. The Hamiltonian is

$$\begin{aligned}
 H = & J \sum_{j=1}^{q-1} \left(\sigma_x^{(j)} \sigma_x^{(j+1)} + \sigma_y^{(j)} \sigma_y^{(j+1)} \right) \\
 & + J\Delta \sum_{j=1}^{q-1} \sigma_z^{(j)} \sigma_z^{(j+1)} + h_z \sum_{j=1}^q \sigma_z^{(j)}.
 \end{aligned} \tag{32}$$

Let R and D denote the single-site raising and lowering matrices. Boundary injection and extraction are described by

$$L_1 = \sqrt{\gamma_1} R^{(1)}, \quad L_2 = \sqrt{\gamma_2} D^{(q)}. \tag{33}$$

The parameters are $J = 0.5$, $\Delta = 0.4$, $h_z = 0.25$, $\gamma_1 = 0.12$, and $\gamma_2 = 0.10$, and the initial state has all spins in the excited product state. For $q = 5, 6, 7$, $T = 1$, $h = 0.1$, and $\tau_r = 10^{-8}$, the simulations give final trace-norm errors 4.4705×10^{-5} , 4.4710×10^{-5} , and 4.4710×10^{-5} , respectively, with final rank 10 in all three cases. This test is in the verified range and almost does not depend on the dimension of Hilbert space, it is controlled by the time step.

Table 7. Large-scale factor storage and indicative single-run wall time at final time 1, step size 0.1, and rank tolerance 10^{-8} .

Spins	Dimension	Rank	Time (s)	Factor MiB	Dense MiB
6	64	10	0.203	0.0098	0.0625
8	256	10	0.663	0.0391	1.0
10	1,024	10	2.015	0.156	16.0
12	4,096	10	4.336	0.625	256
14	16,384	10	9.684	2.500	4,096




Figure 5. Indicative runtime scaling for the boundary-driven spin chain. Small-dimension points include fixed Python and Krylov setup overhead.

Within the range measured, the dimension of the Hilbert space rises by a factor of 256 and the single run time goes from 0.203 s to 9.684 s, roughly a factor of 47.7. The smallest cases are dominated in part by the cost of setting up the Python and Krylov, and so are not used to compute an inferred asymptotic exponent. Even so, they show that the low rank (sparse-operator) implementation can achieve the largest dimension of a tested implementation, without the formation of a dense state/Liouville.

The dimension of the Hilbert space is 16384 for 14 spins. The two are dense density matrix (DDM) and state factor, while state factor uses up to 2.5 MiB while DDM requires 4,096 MiB, implying that state factor's memory consumption is about 1638 times less of the DDM. This reduction is explicitly shown in figure 6. In some sense this figure represents the memory savings on the state representation, and not the memory savings on the overall process, as the calculation still requires memory for sparse Hamiltonian and jump operators, temporary block factors, Krylov bases, etc. The end result remains 10, which is to expected to maintain the delivery of the boundary driven test in the moderate entropy regime focused on by the method.




Figure 6. State-storage comparison for a rank-10 factor and a dense complex128 density matrix. Sparse operators and Krylov workspace are excluded from both curves.

7. Discussion

Five conclusions from the analysis and experiments are: The first is that one can arrange the interaction-picture midpoint rule to exhibit positivity at the algebraic level. Prevents becoming stuck in the state of having a high order approximation that is good in a matrix norm but poor in the state space. Secondly, spectral compression can be measured in the exact same trace norm that is introduced by the global error theorem. Third, the Hermitian residual has a compact indefinite factorization, so that the residual certificate is computable even at a low rank. Fourth, runs of certified variables of step preserve the desired global budgets, with few tries rejected, for both the accumulated residual estimator and the measured error. Fifth, the oscillator rank study shows that there are interesting tradeoffs between time and rank tolerances: If the time tolerance has been exceeded, further gain from ranking is not worth the computational effort.

The certified estimator can be considered understated. It is reliable because it is based on the replacement of the norm of the Bernstein combination with a sum of coefficient norms that are convex. So cancellation between the three Bernstein coefficients is not taken into account. The second effect is part of what the Gauss residual indicator shows and the effect is more material and more sharp. The quadrature indicator can be practically used for step proposals, and the certified Bernstein indicator for acceptance or for the periodic certification. Another way is to partition the interval and use the Bernstein bound on each interval, which will make the estimate more tight, but the proof remains the same.

The present claims are drawn to certain limitations. In the derivation we have assumed an autonomous Lindbladian. Midpoint or commutator-free exponential actions can be used for smooth time-dependent Hamiltonians and jump operators, but the order proof/reconstruction coefficients must then be made time dependent as well. The method is also based on efficient exponential actions. Even if the state is of low rank, that cost could be predominant for dense unstructured generators. Finally, the efficiency of the rank cannot be achieved universally; if the mixing time is long and/or the temperature is high, good dissipation can be obtained up to the Hilbert space dimension, which is the complete rank.

Trace-norm of the projected trace varies linearly with the width of the concatenated factor. If you are going to have many jump operators, the width projected might be the limiting factor. Grouping of operators, random range compression prior to the QR step or tensor-network factorizations are natural solutions. Recent work in tensor-train (TT) and high-order CPTP (CTP) representations [19]–[21] offer interesting lines of future research to integrate the residual formula with even more compact operator representations.

Similar conclusion should also be drawn with respect to the comparison with RK4. It is an example of a lack of structure of a conventional explicit method for a fine step; it is not a claim that RA-SPIF2 beats all the modern Lindblad methods. In other parameter regimes, it could be more desirable to use kraus based high order methods [5] or adaptive variational algorithms[6] or [15]. A unique feature of RA-SPIF2 is the transparency of the second order factor update and the calculation of a posteriori trace-norm certificate.

8. Conclusions

A rank adaptive structure preserving numerical framework for large Lindblad matrix differential equations has been developed. The step of RA-SPIF2 is just an explicit midpoint rule in the interaction picture and decoded into the positive congruence factors. Rank is a quantity controlled by a QR–eigenvalue compression and the discarded spectral mass provides an exact trace-distance perturbation. The accepted factor is positive semidefinite and Construction Normalized.

This a posteriori analysis is based on a quadratic reconstruction of two half steps that are accepted. Exact Lindblad semigroup can be contractive in trace-norm, giving a global residual bound, and exact Lindblad semigroup can be written by Bernstein which gives an exact local estimator. The estimator is in the form of small Hermitian eigenproblems with all the required residual norms involved. By virtue of this property, the certification could be used in the same low rank sparse operating mode the integrator would use.

Numerical results show a close to second order convergence together with proper estimator behavior, explicit adaptive-step control and explicit balance between temporal and rank error. Coarse-step tests also keep positivity to roundoff whereas classical RK4 yields a state that is manifestly negative. The boundary-driven spin-chain calculation is performed at dimension 16384, rank 10 and state factor 2.5MiB. In the future, higher-order residual-certified Kraus schemes and integrating-factor schemes should be developed, the theory should be extended to the time-dependent case for the generator and a combination between the trace-norm estimator and tensor-train state and operator formats should be done.

References

- [1] D. Manzano, “A short introduction to the Lindblad master equation,” *AIP Advances*, vol. 10, no. 2, Art. no. 025106, 2020, doi: 10.1063/1.5115323.
- [2] G. McCauley, B. Cruikshank, D. I. Bondar, and K. Jacobs, “Accurate Lindblad-form master equation for weakly damped quantum systems across all regimes,” *npj Quantum Information*, vol. 6, Art. no. 74, 2020, doi: 10.1038/s41534-020-00299-6.
- [3] A. Trushechkin, “Unified Gorini–Kossakowski–Lindblad–Sudarshan quantum master equation beyond the secular approximation,” *Physical Review A*, vol. 103, Art. no. 062226, 2021, doi: 10.1103/PhysRevA.103.062226.
- [4] Y. Cao and J. Lu, “Structure-preserving numerical schemes for Lindblad equations,” *Journal of Scientific Computing*, vol. 102, Art. no. 27, 2025, doi: 10.1007/s10915-024-02707-x.
- [5] D. Appelö and Y. Cheng, “Kraus is king: High-order completely positive and trace preserving (CPTP) low rank method for the Lindblad master equation,” *Journal of Computational Physics*, vol. 534, Art. no. 114036, 2025, doi: 10.1016/j.jcp.2025.114036.
- [6] L. Gravina and V. Savona, “Adaptive variational low-rank dynamics for open quantum systems,” *Physical Review Research*, vol. 6, Art. no. 023072, 2024, doi: 10.1103/PhysRevResearch.6.023072.

- [7] Y.-T. Chen, C. Farquhar, and R. M. Parrish, “Low-rank density-matrix evolution for noisy quantum circuits,” *npj Quantum Information*, vol. 7, Art. no. 61, 2021, doi: 10.1038/s41534-021-00392-4.
- [8] G. McCaul, K. Jacobs, and D. I. Bondar, “Fast computation of dissipative quantum systems with ensemble rank truncation,” *Physical Review Research*, vol. 3, Art. no. 013017, 2021, doi: 10.1103/PhysRevResearch.3.013017.
- [9] G. Ceruti and C. Lubich, “An unconventional robust integrator for dynamical low-rank approximation,” *BIT Numerical Mathematics*, vol. 62, pp. 23–44, 2022, doi: 10.1007/s10543-021-00873-0.
- [10] G. Ceruti, J. Kusch, and C. Lubich, “A rank-adaptive robust integrator for dynamical low-rank approximation,” *BIT Numerical Mathematics*, vol. 62, pp. 1149–1174, 2022, doi: 10.1007/s10543-021-00907-7.
- [11] M. Hochbruck, M. Neher, and S. Schrammer, “Rank-adaptive dynamical low-rank integrators for first-order and second-order matrix differential equations,” *BIT Numerical Mathematics*, vol. 63, Art. no. 9, 2023, doi: 10.1007/s10543-023-00942-6.
- [12] A. Charous and P. F. J. Lermusiaux, “Stable rank-adaptive dynamically orthogonal Runge–Kutta schemes,” *SIAM Journal on Scientific Computing*, vol. 46, no. 1, 2024, doi: 10.1137/22M1534948.
- [13] F. Nobile and S. Riffaud, “High-order BUG dynamical low-rank integrators based on explicit Runge–Kutta methods,” *Journal of Scientific Computing*, vol. 107, Art. no. 102, 2026, doi: 10.1007/s10915-026-03317-5.
- [14] S. Li and Y. Jiang, “High-order implicit low-rank method with spectral deferred correction for matrix differential equations,” *Journal of Scientific Computing*, vol. 106, Art. no. 75, 2026, doi: 10.1007/s10915-026-03207-w.
- [15] S. Santos, X. Song, and V. Savona, “Low-rank variational quantum algorithm for the dynamics of open quantum systems,” *Quantum*, vol. 9, Art. no. 1620, 2025, doi: 10.22331/q-2025-02-04-1620.
- [16] Z. Ding, X. Li, and L. Lin, “Simulating open quantum systems using Hamiltonian simulations,” *PRX Quantum*, vol. 5, Art. no. 020332, 2024, doi: 10.1103/PRXQuantum.5.020332.
- [17] K. Wang and X. Li, “Simulation-assisted learning of open quantum systems,” *Quantum*, vol. 8, Art. no. 1407, 2024, doi: 10.22331/q-2024-07-11-1407.
- [18] F. Merle and A. Prohl, “A posteriori error analysis and adaptivity for high-dimensional elliptic and parabolic boundary value problems,” *Numerische Mathematik*, vol. 153, pp. 827–884, 2023, doi: 10.1007/s00211-023-01350-2.
- [19] M. Bachmayr, H. Eisenmann, and A. Uschmajew, “Dynamical low-rank tensor approximations to high-dimensional parabolic problems: Existence and convergence of spatial discretizations,” *Numerische Mathematik*, vol. 157, pp. 781–822, 2025, doi: 10.1007/s00211-025-01465-8.
- [20] P. DelMastro, D. Appelö, and Y. Cheng, “Completely positive and trace preserving schemes with tensor train compression for the Lindblad equation,” *arXiv:2605.01494*, 2026, doi: 10.48550/arXiv.2605.01494.
- [21] J. Hu, D. Appelö, and Y. Cheng, “Arbitrary high order low-rank completely positive and trace preserving schemes for Lindblad equations with time-dependent Hamiltonian,” *arXiv:2511.12012*, 2025, doi: 10.48550/arXiv.2511.12012.
- [22] A. Zeynali and Z. Bakhshi, “Adaptive low-rank variational quantum algorithm for simulating dissipative dynamics in photosynthetic complexes,” *Scientific Reports*, 2026, doi: 10.1038/s41598-026-60704-6.