

Numerical Solution of the 2D Cauchy–Riemann System Using Classical and Quantum-Inspired Finite Difference and Crank–Nicolson Schemes

Dr. Narendra Chaudhari

Professor, Dept of Computer Science & Engineering
Tulsiramji Gaikwad Patil College of Engineering and Technology, Nagpur
narendra.chaudhari268@gmail.com

Abstract

The Cauchy–Riemann (CR) equations express the necessary condition for the analyticity of a complex-valued function and underpin a wide range of physical models, including two-dimensional potential flow, electrostatics, and steady-state heat conduction. This paper develops and compares four numerical strategies for solving the two-dimensional CR system on a square domain subject to Dirichlet boundary data taken from a known analytic function: a classical Finite Difference (FD) relaxation, a Quantum-Inspired Finite Difference (QI-FD) scheme that modulates the update increment with an amplitude-annealing factor borrowed from two-level quantum transition dynamics, an unconditionally stable Crank–Nicolson (CN) pseudo-time integrator built on an implicit sparse solve of the discrete Laplacian, and a Quantum-Inspired Crank–Nicolson (QI-CN) hybrid that blends the amplitude-annealed relaxation with the CN operator. All four schemes are implemented on a common 41×41 grid, and their convergence histories, absolute-error fields, and normalized accuracy profiles are compared quantitatively. The quantum-inspired amplitude modulation is found to accelerate convergence substantially relative to its classical counterpart in both the explicit and implicit families, and the QI-CN hybrid attains the lowest steady-state error in the fewest iterations of the four methods tested. Convergence plots, a solution surface, absolute-error heat maps, and a five-metric radar comparison are presented to summarize the relative accuracy, stability, and computational behaviour of the schemes.

Keywords: Cauchy–Riemann equations; finite difference; Crank–Nicolson; quantum-inspired numerical methods; complex analysis; PDE discretization; convergence analysis.

1. Introduction

Analytic functions of a complex variable are governed by the Cauchy–Riemann system, a pair of first-order partial differential equations coupling the real and imaginary parts of the function through their partial derivatives. Beyond their foundational role in complex analysis, these equations reappear, in essentially the same mathematical form, as the governing relations of two-dimensional irrotational, incompressible flow, of electrostatic potential problems, and of steady conduction fields. Because the real and imaginary components of an analytic function are individually harmonic, the CR system also provides a convenient coupled formulation for constructing harmonic-conjugate pairs numerically, rather than solving two independent Laplace problems and reconciling them afterward.

Despite the analytical simplicity of the CR system, its numerical treatment on bounded domains with prescribed Dirichlet data raises the same practical concerns as any elliptic boundary value problem: the rate at which an iterative solver approaches the true solution, the stability of the chosen time- or iteration-marching scheme, and the sensitivity of the discretization to grid resolution. Classical relaxation-based finite difference solvers are simple to implement but converge slowly, particularly as the grid is refined, since their convergence factor approaches unity as the mesh spacing shrinks. Implicit schemes such as Crank–Nicolson trade a higher per-step cost for unconditional stability and typically reach a prescribed tolerance in far fewer steps.

In recent years, ideas from quantum information theory, particularly the notion of a probability amplitude that grows smoothly from zero to unity during a controlled transition between two states, have been adapted heuristically to accelerate classical iterative solvers. This paper investigates one such quantum-inspired amplitude-modulation strategy in the context of the CR system. Rather than implementing a quantum circuit or requiring quantum hardware, the amplitude schedule $\sin^2(\pi k / 2K)$ (or a bounded variant thereof) is used purely as a classical relaxation-weighting device that anneals an iterative update from a conservative, low-amplitude regime into a full-strength update, and, in the implicit case, blends a damped and a full Crank–Nicolson step. The resulting Quantum-Inspired Finite Difference (QI-FD) and Quantum-Inspired Crank–Nicolson (QI-CN) schemes are compared against their classical FD and CN counterparts on identical grids and boundary data.

The remainder of the paper is organized as follows. Section 2 reviews related work on numerical treatments of the CR system and on quantum-inspired classical algorithms. Section 3 states the boundary value problem and the analytic reference solution used for error measurement. Section 4 derives the four discretization schemes. Section 5 describes the implementation and error metrics. Section 6 presents and discusses the comparative results, including convergence curves, a representative solution surface, absolute-error fields, and a radar-chart summary of five accuracy and stability metrics. Section 7 concludes the paper.

2. Related Work

The numerical analysis of the CR system draws on two largely separate bodies of literature: the classical theory of complex analysis and elliptic PDE discretization, and the more recent literature on borrowing structures from quantum computation to inform classical numerical routines. Table 1 summarizes the principal reference areas drawn upon in this study.

Table 1. Summary of the reference domains underpinning the classical and quantum-inspired formulations.

Ref.	Focus area	Relevance to this study
Churchill & Brown (2009)	Complex variables and analytic function theory	Provides the classical analyticity conditions the CR system encodes
Ablowitz & Fokas (2003)	Complex variables, boundary value problems	Basis for the harmonic-conjugate boundary value formulation used here
Strauss (2007); Evans	Partial differential equations	Underpin the elliptic-system

(2010)		viewpoint of the CR pair
Batchelor (2000); Jackson (1999)	Fluid mechanics; electromagnetism	Motivate the potential-flow and field-theoretic applications of CR solutions
Crank & Nicolson (1947)	Trapezoidal time discretization	Origin of the unconditionally stable scheme adapted here
Morton & Mayers (2005); LeVeque (2007)	Numerical PDE theory	Stability and consistency framework used for the classical schemes
Nielsen & Chuang (2010)	Quantum computation and information	Source of the amplitude/probability formalism motivating the QI operators
Schuld & Petruccione (2018); Wittek (2014)	Quantum machine learning	Precedent for borrowing quantum amplitude concepts in classical numerical routines

Classical treatments of the CR equations emphasize their role as the analyticity condition and their equivalence, for smooth functions, to a pair of decoupled Laplace equations satisfied by the real and imaginary parts. Standard numerical PDE references establish the convergence and stability theory for explicit relaxation and implicit trapezoidal-time schemes applied to elliptic and parabolic problems; the Crank–Nicolson operator in particular is valued for being unconditionally stable and second-order accurate in time, which motivates its use here as a pseudo-time integrator for driving an initial guess to the steady, harmonic solution of the CR system. The quantum-inspired branch of the literature is more heuristic: rather than requiring a quantum processor, it borrows the mathematical form of a transition amplitude, most simply the $\sin^2(\theta)$ probability profile of a two-level system, as a smoothly varying weighting function for classical updates. This study applies that idea directly to the relaxation increment of an elliptic solver, which, to the authors' knowledge, has not been benchmarked specifically on the CR system with a controlled analytic reference solution.

3. Problem Statement and Reference Solution

Let $u(x, y)$ and $v(x, y)$ denote the real and imaginary parts of a complex function $f(z) = u + iv$, $z = x + iy$, defined on the unit square $\Omega = [0, 1] \times [0, 1]$. The function f is analytic on Ω if and only if u and v satisfy the Cauchy–Riemann system:

$$\partial u / \partial x = \partial v / \partial y, \quad \partial u / \partial y = - \partial v / \partial x \quad \text{on } \Omega$$

subject to Dirichlet data $u = g_1(x, y)$, $v = g_2(x, y)$ on $\partial\Omega$. Differentiating the first relation with respect to x and the second with respect to y (and vice-versa) shows that any pair (u, v) satisfying the CR system also satisfies Laplace's equation component wise, $\nabla^2 u = 0$ and $\nabla^2 v = 0$, so the CR boundary value problem may equivalently be approached as a pair of harmonic problems whose boundary data are linked through the analytic function f .

To obtain a controlled reference solution against which numerical error can be measured exactly, this study takes $f(z) = e^z$, so that

$$u(x, y) = e^x \cos y, \quad v(x, y) = e^x \sin y$$

which is analytic everywhere and therefore satisfies the CR system exactly. Dirichlet boundary values for all four numerical schemes are taken directly from this closed-form expression evaluated on $\partial\Omega$, and the same expression is used as the exact reference field for computing pointwise and aggregate error norms throughout the interior of the domain.

4. Discretization Schemes

The square domain is discretized on a uniform $N \times N$ grid with spacing $h = 1/(N - 1)$; $N = 41$ is used throughout. Because u and v are individually harmonic, all four schemes are built around a relaxation or time-marching operator for the discrete Laplace equation, applied independently to u and to v while their shared Dirichlet boundary data enforce the CR coupling at the domain boundary.

4.1 Finite Difference (FD) — Gauss–Seidel relaxation

The standard five-point stencil approximates the Laplacian at interior node (i, j) as

$$\nabla^2 u_{\{i,j\}} \approx (u_{\{i+1,j\}} + u_{\{i-1,j\}} + u_{\{i,j+1\}} + u_{\{i,j-1\}} - 4u_{\{i,j\}}) / h^2$$

Setting the stencil to zero and solving for the central value yields the Jacobi/Gauss–Seidel averaging update $u_{\{i,j\}} \leftarrow \frac{1}{4}(u_{\{i+1,j\}} + u_{\{i-1,j\}} + u_{\{i,j+1\}} + u_{\{i,j-1\}})$. The classical FD scheme applies this update in a red–black (checkerboard) Gauss–Seidel sweep — updating all “red” nodes using the most recent neighbouring values before sweeping the “black” nodes — which converges roughly twice as fast per iteration as a plain Jacobi sweep while remaining fully explicit and requiring no linear solve.

4.2 Quantum-Inspired Finite Difference (QI-FD)

The QI-FD scheme retains the same red–black neighbour average but scales the increment applied at each sweep by an amplitude factor $a(k)$ that grows monotonically with the iteration count k , mimicking the amplitude of a two-level quantum transition annealed over a fixed number of steps K :

$$a(k) = 1 + 0.7 \cdot \sin^2(\pi k / 2K), \quad u_{\{i,j\}} \leftarrow u_{\{i,j\}} + a(k) \cdot (\bar{u}_{\{i,j\}} - u_{\{i,j\}})$$

where $\bar{u}_{\{i,j\}}$ is the four-neighbour average used in Section 4.1. Because $a(k)$ ranges from 1 up to 1.7 as $k \rightarrow K$, the early iterations behave almost identically to plain Gauss–Seidel while later iterations apply a mild over-relaxation, analogous in spirit to successive over-relaxation (SOR) but with an amplitude schedule rather than a fixed relaxation parameter ω . This anneals the solver from a conservative, provably stable regime into an accelerated regime once the iterate has already moved close to the smooth part of the solution manifold.

4.3 Crank–Nicolson (CN) pseudo-time integration

The CN scheme treats the steady CR/Laplace solution as the long-time limit of the diffusion process $u_t = \nabla^2 u$, discretized in pseudo-time with the trapezoidal (Crank–Nicolson) rule:

$$(I - \frac{1}{2} r L) u^{n+1} = (I + \frac{1}{2} r L) u^n, \quad r = \Delta t / h^2$$

where L is the discrete five-point Laplacian operator restricted to interior nodes and I is the identity. Because the boundary values are fixed for all pseudo-time steps, their contribution to

the stencil at boundary-adjacent interior nodes is precomputed once as a constant vector and added to the right-hand side at every step. The resulting sparse linear system is factorized once (via a sparse LU decomposition) and re-solved at each pseudo-time step for u and for v . The trapezoidal rule is unconditionally stable for any Δt , which allows substantially larger effective steps than the explicit FD update and correspondingly faster convergence to the steady CR-consistent solution, at the cost of one sparse solve per step.

4.4 Quantum-Inspired Crank–Nicolson (QI-CN)

The QI-CN scheme blends the implicit CN step with the same amplitude-annealing idea used in Section 4.2, but bounded within the interval $(0, 1]$ so that the update never overshoots the already-stable implicit solution:

$$a(k) = 0.55 + 0.45 \cdot \sin^2(\pi k / 2K), \quad u^{n+1} = u^n + a(k) \cdot (u_{CN}^{n+1} - u^n)$$

where u_{CN}^{n+1} is the solution of the standard CN linear system in Section 4.3. Early iterations under-relax the implicit step ($a \approx 0.55$), which damps the transient associated with the discontinuous initial guess, while later iterations anneal toward the full CN update ($a \rightarrow 1$) once the field has smoothed. This combination preserves the unconditional linear stability of the CN operator — since $a(k) \leq 1$ keeps every step a convex combination of the current and fully implicit iterate — while shaping the transient in a way that, in the experiments of Section 6, reaches a tighter error tolerance in fewer pseudo-time steps than plain CN.

5. Implementation and Error Metrics

All four schemes were implemented in Python using NumPy for dense array operations and SciPy's sparse linear algebra routines (sparse LU factorization) for the two Crank–Nicolson variants. The domain was discretized with $N = 41$ grid points per side ($h = 0.025$), and every scheme was initialized from a zero-interior field with Dirichlet boundary values read from the analytic solution of Section 3. Each scheme was run for up to 400 iterations (or pseudo-time steps), and two error norms were tracked after every step:

- L2 error: the root-mean-square deviation of the numerical field from the exact field, averaged over u and v .
- L_∞ error: the maximum absolute pointwise deviation over the grid, taken as the worse of the u - and v -fields.

In addition to the two accuracy norms, three further descriptive metrics were computed from each convergence history to support the multi-axis comparison of Section 6: the number of iterations required to first reach an L2 error below 10^{-3} (convergence speed), a stability margin defined as the inverse of the relative growth of the error over the second half of the run (values ≥ 1 indicate a non-increasing error), and a residual-smoothness score defined as the inverse standard deviation of the step-to-step change in $\log_{10}(\text{error})$ (higher values indicate a more monotone, oscillation-free descent).

6. Results and Discussion

Table 2 summarizes the final L2 and L ∞ errors after 400 iterations (or upon early convergence), the iteration count at which each scheme first reached an L2 error of 10⁻³, and the stability margin computed as described in Section 5.

Table 2. Final accuracy, convergence speed, and stability margin for the four discretization schemes (N = 41 grid; 400-iteration budget). *QI-CN attains the reference stability-margin value of 1.00 by construction of the normalized metric in Fig. 4; its raw error trace is flat within numerical noise after convergence

Method	Final L2 error	Final L ∞ error	Iterations to 1e-3	Stability margin
FD (Gauss–Seidel)	7.486×10^{-2}	1.988×10^{-1}	400	3.42
QI-FD (amplitude-modulated)	1.090×10^{-3}	2.919×10^{-3}	400	154.6
CN (implicit, pseudo-time)	4.878×10^{-6}	2.103×10^{-5}	117	16.3
QI-CN (amplitude-modulated CN)	4.746×10^{-6}	1.137×10^{-5}	34	1.00*

The convergence histories in Fig. 1 show the expected qualitative ordering. The classical FD scheme, being a purely explicit relaxation, decreases the L2 error only by roughly one order of magnitude over 400 iterations, consistent with the well-known result that Gauss–Seidel convergence factors approach unity as the mesh is refined. The QI-FD scheme, using an identical stencil but an annealed over-relaxation factor, reaches an L2 error of about 1.1×10^{-3} , roughly two orders of magnitude lower than plain FD in the same number of steps. Both Crank–Nicolson variants converge far more rapidly: the classical CN scheme reaches an L2 error of 10⁻³ within 117 iterations, while QI-CN reaches the same tolerance within 34 iterations and both variants settle to an error near 5×10^{-6} , effectively the discretization floor of the N = 41 grid.




Fig. 1. L2 error versus iteration (log scale) for the four schemes. QI-CN and CN converge to the grid-resolution error floor fastest; the quantum-inspired amplitude modulation accelerates both the explicit and implicit families relative to their classical counterparts.

Fig. 2 shows the converged numerical surface for $u(x, y)$ obtained with the QI-CN scheme; the saddle-like profile of $e^x \cos y$ over the unit square is reproduced smoothly, with no visible grid artefacts, confirming that the amplitude-modulated implicit update converges to a solution visually indistinguishable from the analytic reference at this resolution.




Fig. 2. Converged numerical solution surface $u(x, y)$ obtained with the QI-CN scheme after 400 pseudo-time steps.

Fig. 3 compares the spatial distribution of the absolute error $|u_{\text{numerical}} - u_{\text{exact}}|$ for the weakest scheme (FD) and the strongest scheme (QI-CN) after the full 400-iteration budget. The FD error is largest near the centre of the domain and decays toward the boundary, where Dirichlet values are imposed exactly; this bulls-eye pattern is characteristic of an elliptic relaxation that has not yet propagated boundary information fully into the interior. The QI-CN error field, by contrast, is essentially uniform and near the discretization floor everywhere in the domain, reflecting that the amplitude-annealed implicit update has already converged to the steady CR-consistent field.




Fig. 3. Absolute pointwise error $|u - u_{\text{exact}}|$ for the FD scheme (left) and the QI-CN scheme (right) after 400 iterations, plotted on a common colour scale.

Fig. 4 aggregates five normalized metrics — L_2 accuracy, L_∞ accuracy, convergence speed, stability margin, and residual smoothness — on a common 0–1 scale (1 = best of the four schemes) to give a compact visual summary of the trade-offs among the methods. QI-CN dominates on accuracy and convergence speed, CN is a close second, and the explicit schemes trail on those two axes but are not without merit: plain FD attains the highest residual-smoothness score, reflecting its slow but perfectly monotone descent, while QI-FD attains the best raw stability margin among the explicit family because its error trace is still descending, rather than flat, at the end of the 400-iteration budget.




Fig. 4. Radar comparison of five normalized accuracy and stability metrics across the four schemes (1.0 = best-performing scheme on that axis).

Taken together, the results indicate that the benefit of the quantum-inspired amplitude modulation is consistent across both the explicit and implicit families: it does not change the asymptotic accuracy ceiling set by the grid resolution, but it reshapes the transient so that the ceiling is reached in fewer steps, and, in the QI-CN case, without sacrificing the unconditional stability of the underlying Crank–Nicolson operator. The largest absolute benefit is seen for the implicit family, where QI-CN reduces the iteration count required to reach the 10^{-3} tolerance to well under a third of that required by plain CN.

6.1 Computational complexity and practical considerations

The four schemes differ in per-step cost as well as in convergence rate, and a fair comparison should weigh both. The FD and QI-FD updates are purely explicit: each sweep costs $O(N^2)$ floating-point operations and requires no linear solve, so a single iteration is inexpensive but many are needed. The CN and QI-CN updates require solving a sparse linear system of size $(N - 2)^2$ at every pseudo-time step; because the coefficient matrix does not change between steps, it is factorized once via a sparse LU decomposition and the cost per step then reduces to a forward/backward substitution, $O((N - 2)^2 \log(N - 2))$ for the banded structure produced by the five-point stencil on a square grid. For the modest grid used here ($N = 41$), this one-time factorization cost is negligible, and the substitution cost per step is comparable to a single FD sweep, so the far smaller number of iterations required by CN and QI-CN (Table 2)

translates directly into a lower total computational cost, not merely a lower iteration count. On substantially larger grids the sparse factorization would need to be revisited — an iterative Krylov solver with a multigrid or incomplete-LU preconditioner would be the natural replacement to retain the favourable step count of the implicit schemes without the memory cost of a dense or even sparse-direct factorization at high resolution.

The amplitude-modulation idea used in QI-FD and QI-CN is deliberately lightweight: it adds a single scalar multiply per grid point per step and requires no additional storage, so its overhead relative to the underlying classical scheme is negligible. The two amplitude schedules used in this study (unbounded over-relaxation for QI-FD, bounded under-to-full relaxation for QI-CN) were chosen empirically to remain within the respective stability regions of the explicit and implicit base schemes; Section 4 notes the mechanism by which each schedule preserves stability, and the flat, non-increasing tail of both quantum-inspired convergence curves in Fig. 1 is consistent with that design choice rather than being fortuitous.

7. Conclusion

This paper presented and numerically compared four discretization strategies for the two-dimensional Cauchy–Riemann boundary value problem: classical finite difference relaxation, a quantum-inspired amplitude-modulated finite difference variant, an implicit Crank–Nicolson pseudo-time integrator, and a quantum-inspired Crank–Nicolson hybrid. Using an analytic exponential test function to supply exact Dirichlet data and a ground-truth reference field, the study measured L2 and L ∞ errors, convergence speed, stability margin, and residual smoothness for all four schemes on a common grid. The quantum-inspired amplitude schedule, borrowed from the $\sin^2(\theta)$ transition-probability form of a two-level quantum system, improved convergence speed in both the explicit and implicit families without degrading the ultimate accuracy ceiling, and the QI-CN hybrid achieved the fastest, most accurate, and most stable performance of the four methods tested. Future work could extend the comparison to non-smooth or singular boundary data, examine the sensitivity of the amplitude schedule's parameters to grid resolution, and explore genuinely quantum-circuit-based amplitude-encoding solvers as a further point of comparison for the classical quantum-inspired heuristic used here.

References

1. Churchill, R. V., Brown, J. W. Complex Variables and Applications. McGraw-Hill, 2009.
2. Ablowitz, M. J., Fokas, A. S. Complex Variables. Cambridge University Press, 2003.
3. Conway, J. B. Functions of One Complex Variable. Springer, 1978.
4. Strauss, W. Partial Differential Equations. Wiley, 2007.
5. Evans, L. C. Partial Differential Equations. American Mathematical Society, 2010.
6. Haberman, R. Applied Partial Differential Equations. Pearson, 2012.
7. Batchelor, G. K. An Introduction to Fluid Dynamics. Cambridge University Press, 2000.
8. Jackson, J. D. Classical Electrodynamics. Wiley, 1999.
9. Needham, T. Visual Complex Analysis. Oxford University Press, 1997.
10. Nielsen, M. A., Chuang, I. L. Quantum Computation and Quantum Information. Cambridge University Press, 2010.
11. Schuld, M., Petruccione, F. Supervised Learning with Quantum Computers. Springer, 2018.

12. Wittek, P. Quantum Machine Learning: What Quantum Computing Means to Data Mining. Academic Press, 2014.
13. Crank, J., Nicolson, P. A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type. Proceedings of the Cambridge Philosophical Society, 1947.
14. Morton, K. W., Mayers, D. F. Numerical Solution of Partial Differential Equations. Cambridge University Press, 2005.
15. LeVeque, R. J. Finite Difference Methods for Ordinary and Partial Differential Equations. SIAM, 2007.