ORIGINAL RESEARCH ARTICLE | UHD JOURNAL OF SCIENCE AND TECHNOLOGY

Spectral Analysis and Long-Term Behavior of Discrete-Time Markov Chains

Abdul Salam Hemat[1], Abdullah Abid[2], Ahmad Ali Mehrad[1]

[1]Department of Mathematics, Faculty of Education, Ghazni University, Ghazni, Afghanistan, [2]Department of Mathematics, Faculty of Education, Shikh Zayed University, Khost, Afghanistan

ABSTRACT

This paper develops a spectral-probabilistic framework for examining the long-term dynamics of finite discrete-time Markov chains (DTMCs). The framework combines stochastic interpretation with eigenvalue-eigenvector analysis, explicit structural validation, probability propagation, and convergence diagnostics implemented in MATLAB. The principal four-state benchmark has transition matrix P and initial distribution $x_0 = [1, 0, 0, 0]$. For the empirical tolerance $\varepsilon = 10^{-3}$, the benchmark trajectory reaches stationarity within 13 steps, with stationary distribution $\pi = (0.2313, 0.2916, 0.2676, 0.2094)$, subdominant spectral radius $\rho = 0.5911$, and spectral gap $\gamma = 0.4089$. The 13-step result is explicitly tied to the stated initial condition and tolerance. A formal worst-case calculation over all initial distributions gives the same 13-step threshold for this benchmark at $\varepsilon = 10^{-3}$. A conservative spectral-gap reference scale is about 21 steps; because the benchmark is not assumed reversible, this value is reported as a comparison scale rather than as a direct theorem for the matrix. Additional synthetic matrices representing slow mixing, periodicity, and near-decomposability demonstrate that the diagnostic procedure distinguishes structurally different regimes. The study is therefore presented as a synthetic numerical benchmark rather than as a validated real-world reliability, deterioration, or risk model.

Index Terms: Discrete-Time Markov Chains, Spectral Analysis, Stationary Distribution, Convergence Rate, Spectral Gap, Eigenvalues, Periodicity, Mixing Time, MATLAB

1. INTRODUCTION

Discrete-time Markov chains (DTMCs) constitute one of the fundamental frameworks in stochastic modeling, providing a rigorous basis for describing systems that evolve through probabilistic transitions over discrete time steps. Their flexibility has supported research on crop-yield prediction, market-switching behavior, portfolio optimization, satellite-constellation spare-strategy analysis, Markovian jump-system stability, reinforcement learning, communication networks, climate-related modeling, digital twins, open networks, periodically observed stochastic processes, and other stochastic systems [1]–[10], [12], [13]. In these settings, an important question is not only how a system moves between states, but also how its long-term behavior develops, whether it approaches equilibrium, and how rapidly that approach occurs.

Traditional approaches to DTMC analysis include iterative state-probability propagation, simulation, and steady-state estimation [1], [4], [11], [17]. These procedures are useful for numerical approximation and scenario testing, but state trajectories alone do not always reveal the structural mechanisms associated with convergence, oscillation, slow modes, or nearly decomposable behavior. Spectral analysis complements direct iteration by connecting these observable trajectories to the operator structure of the transition matrix.

Access this article online

DOI:10.21928/uhdjst.v10n2y2026.pp188-197 | E-ISSN: 2521-4217 | P-ISSN: 2521-4209

Copyright © 2026 Hemat, *et al*. This is an open access article distributed under the Creative Commons Attribution Non-Commercial No Derivatives License 4.0 (CC BY-NC-ND 4.0)

Corresponding author's e-mail: Abdul Salam Hemat, Department of Mathematics, Faculty of Education, Ghazni University, Ghazni, Afghanistan. E-mail: abdulsalamh616@gmail.com

Received: 01-08-2026 Accepted: 30-09-2026 Published: 14-10-2026

A more informative perspective is obtained by examining the eigenvalues and eigenvectors of the transition matrix. Spectral quantities describe invariant behavior and the decay of transient components, while eigenvalues near the unit circle can indicate slow relaxation or oscillatory effects [14], [15], [18], [19]. The present study uses these ideas to construct a computational pipeline that combines structural checks with spectral and probability-based diagnostics.

The objective is to provide a reproducible computational procedure that: (i) Validates the transition matrix; (ii) checks irreducibility explicitly; (iii) identifies periodicity using graph structure; (iv) computes the stationary distribution and subdominant spectral radius; (v) measures convergence from a stated initial condition; and (vi) evaluates the formal worst-case total-variation distance over all initial distributions. The numerical study is deliberately framed as a synthetic benchmark. Claims concerning reliability engineering, deterioration modeling, or risk assessment are therefore treated as potential application domains rather than as demonstrated applications in this paper.

2. THEORETICAL AND COMPUTATIONAL SPECTRAL ANALYSIS OF LONG-TERM BEHAVIOR IN DISCRETE-TIME MARKOV CHAINS

Let P be the transition matrix of a finite DTMC and let x_t denote a row probability vector. The temporal evolution is

$$x_{t+1} = x_t P$$

A stationary distribution π satisfies

$$\pi P = \pi, \sum_i \pi_i = 1, \pi_i \geq 0$$

Under irreducibility, a finite chain has a unique stationary distribution. If the chain is also aperiodic, the state distribution converges to π from every initial distribution. The dominant eigenvalue is $\lambda_1 = 1$, and the magnitude of the largest non-trivial eigenvalue is denoted by

$$\rho = max\{|\lambda| : \lambda \in \sigma(P), \lambda \neq 1\}, \gamma = 1 - \rho$$

The quantity γ is the spectral gap. A smaller ρ, equivalently a larger γ, is associated with faster asymptotic decay of non-stationary spectral components. This relationship is informative but does not by itself replace direct mixing-time calculations, particularly for non-reversible chains.

For a probability vector x_t, convergence is measured using total variation distance and the Euclidean norm:

$$TV(x_t,\pi) = ½\|x_t - \pi\|_1, \|x_t - \pi\|_2 = [\sum_i (x_{t,i} - \pi_i)^2]^{1/2}$$

The empirical mixing time used in this study is defined at a stated tolerance ε by

$$t_{mix}(\varepsilon) = min\{t \geq 0 : max_x \, TV(xP^t, \pi) \leq \varepsilon\}$$

For a finite state space, the maximum over the probability simplex is attained at an extreme point. Consequently, the worst-case calculation can be implemented by evaluating the rows of P^t and taking the maximum of their total-variation distances from π. This is stronger than reporting convergence from a single initial state.

Spectral-gap behavior in this study is characterized by the directly computed quantities ρ and γ, together with the empirical mixing time at $\varepsilon = 10^{-3}$ and the worst-case mixing time. This reporting keeps the interpretation tied to computed diagnostics rather than to arbitrary numerical classifications.

3. BACKGROUND AND RELATED WORK

The asymptotic analysis of DTMCs is traditionally grounded in recurrence, irreducibility, periodicity, and stationary distributions. Spectral methods complement this probabilistic framework by relating invariant behavior and transient decay to the transition operator [14], [15], [18].

Recent literature represented by Refs. [14]–[19] spans several related but distinct directions, including reversible Markov-chain spectral theory, stochastic matrix factorization, mixing and dynamical-systems results, hidden Markov models, classical Markov-chain limit theory, and spectral analysis of diffusion models. Table 1 reports these references according to the subjects explicitly represented by their bibliography entries, avoiding mismatches between a reference number and an unrelated research description.

The contribution of the present paper is methodological: It combines matrix validation, graph-based irreducibility and periodicity checks, spectral diagnostics, empirical mixing-time calculations, and visual reporting in one MATLAB workflow. The additional numerical experiments are used to test the diagnostic behavior across structurally different synthetic chains.

TABLE 1. Background and related work on the long-term analysis of discrete-time Markov chains

References	Documented topic from bibliography	Role in the present study
[14]	Spectral analysis of reversible Markov chains; finite-time spectral properties.	Supports spectral interpretation of Markov-chain eigenstructure and provides related reversible-chain context.
[15]	Spectral theory for Markov chains with transition matrices admitting stochastic bidiagonal factorization.	Supports the operator-based spectral treatment of stochastic transition matrices.
[16]	Statistical limit theorems for Axiom A diffeomorphisms and related dynamical systems.	Provides related dynamical-systems context for mixing and limit behavior; it is not used as a direct DTMC-specific source.
[17]	Extending conformal prediction to hidden Markov models.	Provides related hidden-state Markov-model context; it is not used as the theoretical source for periodicity detection.
[18]	Limit theorems for Markov chains and stochastic properties of dynamical systems.	Provides classical theoretical context for long-term and limit behavior of Markov chains.
[19]	Spectral analysis of diffusion models with application to schedule design.	Provides a modern related example of spectral analysis outside the present synthetic DTMC benchmark.

DTMC: Discrete-time Markov chains

4. METHODOLOGY AND MATLAB IMPLEMENTATION

The computational workflow contains five structural components: Stochastic validation, irreducibility verification, periodicity detection, spectral analysis, and convergence assessment.

The transition matrix is first checked for square dimensions, non-negative entries, and unit row sums. The initial distribution is checked for non-negativity and unit sum.

Second, irreducibility is verified computationally from the directed transition graph. The implementation constructs the Boolean reachability relation and applies a transitive-closure procedure. The chain is declared irreducible only when every state can reach every other state.

Third, periodicity is determined from graph cycle structure. The algorithm performs a breadth-first search from one state, records graph distances, and computes the greatest common divisor of the quantities d(i)+1-d(j) over directed edges i > j. For an irreducible chain, this gives the period. In particular, a period-2 chain with eigenvalue -1 is detected even though the eigenvalue is purely real.

Fourth, the program computes the eigenvalues and eigenvectors of P^t, obtains the stationary distribution from the eigenvector associated with λ approximately equal to 1, and computes ρ and γ. Spectral quantities are reported directly, without converting γ into qualitative classes using arbitrary numerical cutoffs.

Fifth, probability propagation is performed from x_0 = [1, 0, 0, 0]. Total variation and L_2 distances are recorded. In parallel, P^t is generated for every time step, and the largest row-wise total-variation distance from π is recorded. This provides the formal worst-case mixing diagnostic over all initial distributions.

t_{mix} (10^{-3}; x_0) = 13 steps, t_{mix} (10^{-3}; worst case) = 13 steps

The benchmark also permits a conservative spectral-gap reference calculation. Using $\varepsilon = 10^{-3}$, π_min = min_i π_i approximately 0.2094, and $\gamma = 0.4089$ gives the reference scale

$$t_{gap} \approx \frac{log[\frac{1}{\varepsilon \pi_{min}}]}{\gamma} = 21\, steps$$

This 21-step quantity is included only as a standard spectral-gap comparison scale. The benchmark is not assumed to satisfy reversibility, so the value is not presented as a direct theorem guaranteeing the observed mixing time. The empirical worst-case calculation is the formal quantity used for the conclusions.

The complete MATLAB implementation used for the numerical analysis is provided in Appendix A.

5. RESULTS AND DISCUSSION

The benchmark transition matrix is

$$\begin{bmatrix} 0.65 & 0.2 & 0.1 & 0.05 \\ 0.15 & 0.55 & 0.2 & 0.1 \\ 0.1 & 0.2 & 0.5 & 0.2 \\ 0.05 & 0.15 & 0.25 & 0.55 \end{bmatrix}$$

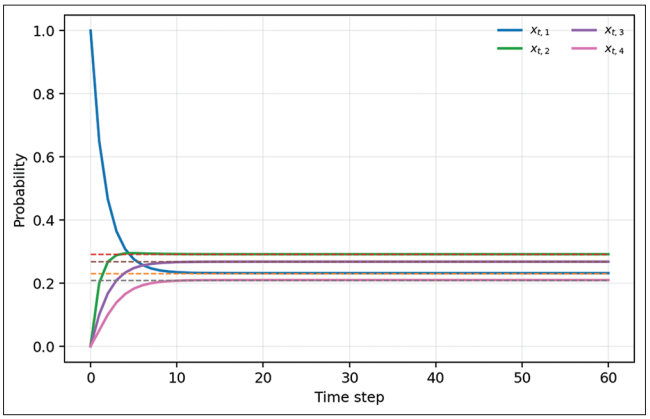

Fig. 1. State-probability evolution from x_0 = [1, 0, 0, 0] toward the stationary distribution.

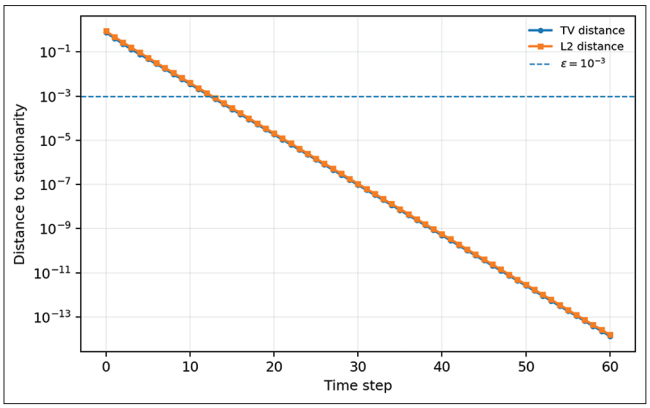

Fig. 2. Total-variation and L_2 convergence diagnostics for the benchmark chain.

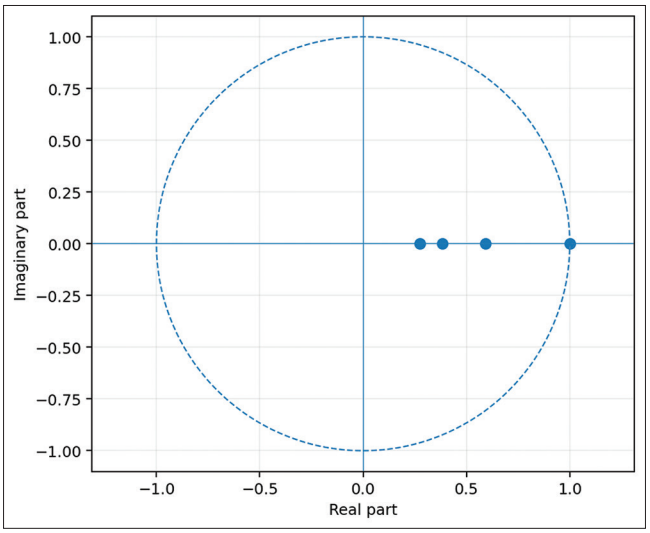

Fig. 3. Eigenvalue spectrum of the benchmark chain with the unit-circle boundary.

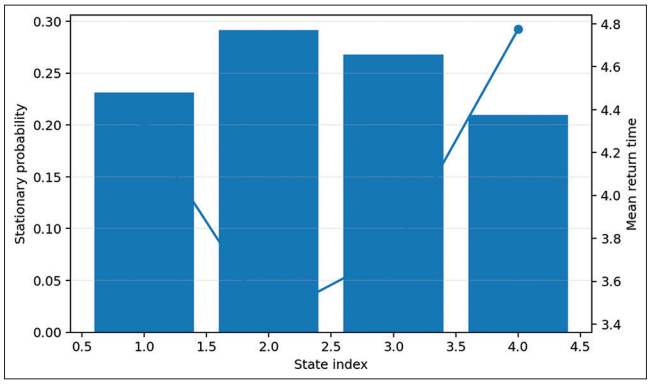

Fig. 4. Stationary probabilities and mean return times for the four benchmark states.

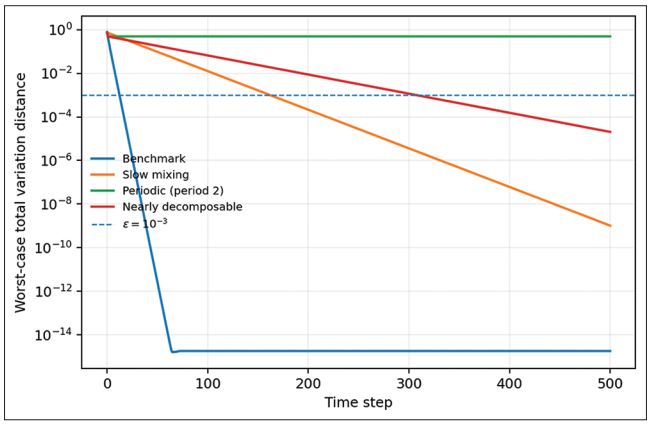

Fig. 5. Worst-case total-variation convergence across the four structural regimes.

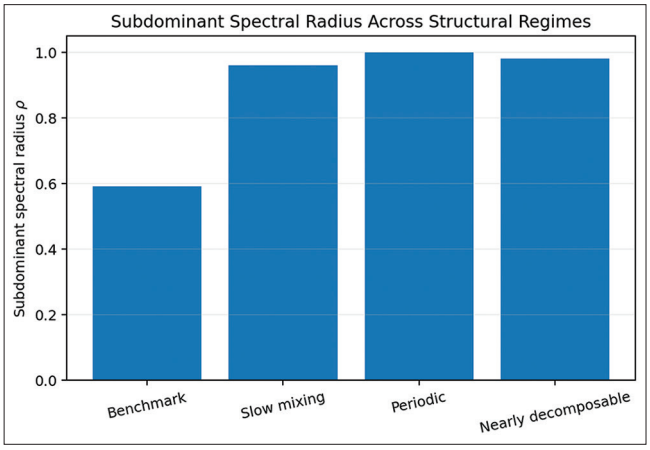

Fig. 6. Subdominant spectral radius for the benchmark and additional structural regimes.

The initial condition is x_0 = [1, 0, 0, 0]. The computed stationary distribution is π = [0.231348, 0.291581, 0.267625,

0.209446]. The eigenvalues are approximately 1.0000, 0.5911, 0.3844, and 0.2745. Hence $\rho = 0.5911$ and $\gamma = 0.4089$. The graph-based irreducibility check returns true, and the graph-based period is 1.

At $t = 0$, the total-variation distance is TV = 0.7687. At $t = 1$, the value is approximately 0.4187. Using $\varepsilon = 10^{-3}$, the trajectory first crosses the threshold at $t = 13$, where TV is approximately 7.29×10^{-4}.

TABLE 2. Computational workflow

Step	Procedure	Main output
1	Define finite state space and transition matrix	Stochastic model
2	Validate dimensions, non-negativity, and row sums	Valid row-stochastic matrix
3	Check graph reachability by transitive closure	Irreducibility flag
4	Compute graph period by cycle-length gcd	Period and periodicity flag.
5	Compute eigenvalues/ eigenvectors of P^t	Spectral profile and stationary vector
6	Compute rho and gamma without arbitrary classification thresholds	Subdominant spectral radius and gap
7	Propagate x0 and calculate TV/L2 distances	Empirical convergence trajectory
8	Evaluate row-wise TV distances of P^t	Formal worst-case mixing time
9	Compute stationary probabilities and 1/pi_i	Long-run occupancy and mean return times
10	Generate figures and summary diagnostics	Reproducible numerical report

TABLE 3. Diagnostic comparison at $\varepsilon=10^{-3}$

Chain	rho	Gamma	Period	Worst-case mixing time
Benchmark	0.5911	0.4089	1	13
Slow mixing	0.9600	0.0400	1	163
Periodic (period 2)	1.0000	0.0000	2	No finite mixing time
Nearly decomposable	0.9800	0.0200	1	308

The state-probability evolution and convergence diagnostics are shown in Figs. 1 and 2, respectively, while the eigenvalue spectrum is presented in Fig. 3. The stationary probabilities and mean return times are shown in Fig. 4; the worst-case total-variation convergence and subdominant spectral radii across the structural regimes are compared in Figs. 5 and 6, respectively.

The complete computational workflow is summarized in Table 2.

By $t = 60$, the benchmark distances are at approximately the numerical precision level, with TV on the order of 10^{-14}. The semilogarithmic trajectory is consistent with geometric decay associated with the subdominant spectral radius.

The spectrum contains four real eigenvalues and no non-trivial eigenvalue on the unit circle. This agrees with the graph-based period calculation of 1 and supports convergence of the benchmark chain. The spectral result is used as a diagnostic explanation of the observed trajectory, not as a substitute for direct convergence measurement.

The mean return times $1/\pi_i$ are approximately 4.3225, 3.4296, 3.7366, and 4.7745. The state with the largest stationary probability therefore has the smallest mean return time, consistent with the standard recurrence relation for finite irreducible positive recurrent chains.

To assess convergence beyond a single initial condition, the analysis evaluates the maximum TV distance over all initial distributions. Because $TV(xP^t, \pi)$ is convex in x, the maximum over the probability simplex occurs at an extreme point; hence it is sufficient to examine the rows of P^t. For the benchmark, the maximum row-wise TV distance at $t = 13$ is 7.29×10^{-4}, so the formal worst-case mixing time at $\varepsilon = 10^{-3}$ is also 13 steps. The initial-condition result and the worst-case result are therefore consistent for this matrix.

TABLE 4. Summary of the numerical study

Diagnostic	Benchmark	Additional regimes	Interpretation
Irreducibility	Verified computationally	Period-2 example also irreducible	Graph reachability is explicitly tested.
Periodicity	Period 1	Period-2 example detected	Graph cycle analysis detects lambda = −1.
Spectral radius	rho=0.5911	0.96, 1.00, 0.98	Higher rho corresponds to slower spectral decay in these examples.
Worst-case TV mixing	13 steps at epsilon=10^{-3}	163, no finite mixing time for period-2, and 308 steps	Finite-time behavior is directly measured.
Application scope	Synthetic benchmark	Synthetic structural tests	Real-world application remains future work.

$$max_i\ TV\ (e_i P^{13}, \pi) = 7.29 \times 10^{-4} < 10^{-3}$$

The conservative spectral-gap reference scale is approximately 21 steps. The empirical threshold is earlier than this comparison scale, which is consistent with the conservative nature of such spectral reference bounds and with the fact that the benchmark is not assumed to be reversible.

The principal numerical findings for the benchmark and the additional structural regimes are summarized in Table 4.

5.1. Numerical Diversity across Structural Regimes

To test whether the diagnostic pipeline responds to structural changes rather than only to the benchmark, three additional synthetic transition matrices were analyzed: A slow-mixing chain, an irreducible period-2 chain, and a nearly decomposable chain. All are four-state stochastic matrices so that the comparison remains controlled.

The additional matrices are explicitly defined below.

$$\begin{bmatrix} 0.97 & 0.01 & 0.01 & 0.01 \\ 0.01 & 0.97 & 0.01 & 0.01 \\ 0.01 & 0.01 & 0.97 & 0.01 \\ 0.01 & 0.01 & 0.01 & 0.97 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0.5 & 0 & 0.5 \\ 0.5 & 0 & 0.5 & 0 \\ 0 & 0.5 & 0 & 0.5 \\ 0.5 & 0 & 0.5 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0.49 & 0.5 & 0.005 & 0.005 \\ 0.5 & 0.49 & 0.005 & 0.005 \\ 0.005 & 0.005 & 0.49 & 0.5 \\ 0.005 & 0.005 & 0.5 & 0.49 \end{bmatrix}$$

For the slow-mixing matrix, $\rho = 0.96$ and the stationary distribution is uniform. For the periodic matrix, the eigenvalues include 1 and −1 and the graph period is 2. For the nearly decomposable matrix, $\rho = 0.98$, with two weakly coupled state clusters and a uniform stationary distribution.

The slow-mixing matrix has $\rho = 0.96$ and reaches $\varepsilon = 10^{-3}$ at 163 steps. The nearly decomposable matrix has $\rho = 0.98$ and reaches the threshold at 308 steps. The period-2 matrix has eigenvalues, including −1, has graph period 2, and does not converge to the stationary distribution in total variation; its worst-case TV distance remains bounded away from zero. These cases demonstrate that periodicity cannot be inferred from the imaginary parts of eigenvalues alone and that subdominant spectral radius and graph structure should be examined together.

6. COMPARATIVE ANALYSIS WITH RELATED STUDIES AND SIGNIFICANCE OF FINDINGS

The comparison focuses on what the numerical experiments demonstrate. The benchmark exhibits a clear separation between the unit eigenvalue and the non-trivial spectrum, while the additional matrices show that slow relaxation, periodicity, and near-decomposability produce substantially different diagnostics. This supports the use of a combined spectral-probabilistic workflow rather than reliance on one numerical indicator.

With respect to convergence theory, the benchmark's approximately geometric decay is consistent with the general role of the subdominant spectral component in finite-state Markov dynamics [14], [15], [18]. However, the present study does not claim that one eigenvalue universally determines finite-time mixing for every non-reversible chain. The empirical TV curve and the formal worst-case calculation are therefore reported alongside the spectral quantities.

The additional experiments provide a direct scope check. A chain with $\rho = 0.96$ mixes much more slowly than the benchmark, whereas a nearly decomposable chain with $\rho = 0.98$ exhibits an even longer transient regime. The period-2 example shows that a chain may possess a stationary distribution while failing to converge to it from arbitrary initial conditions because of periodic oscillation. These results illustrate the diagnostic purpose of the framework without asserting that the synthetic matrices represent particular physical systems.

The numerical evidence in this study is based entirely on synthetic transition matrices. Reliability engineering, deterioration modeling, and risk assessment are reserved as potential future application domains for studies using calibrated real-world transition matrices.

7. CONCLUSION AND FUTURE WORK

This study developed an integrated spectral-probabilistic framework for characterizing asymptotic and transient behavior in finite discrete-time Markov chains. The

implementation combines stochastic validation, explicit graph-based irreducibility verification, graph-based periodicity detection, spectral analysis, probability propagation, and empirical mixing diagnostics.

For the four-state benchmark with $x_0 = [1, 0, 0, 0]$ and $\varepsilon = 10^{-3}$, convergence to the stationary distribution occurs at 13 steps. The stationary distribution is $\pi = [0.2313, 0.2916, 0.2676, 0.2094]$, the subdominant spectral radius is $\rho = 0.5911$, and the spectral gap is $\gamma = 0.4089$. The TV distance is 0.7687 at $t = 0$ and approximately 0.4187 at $t = 1$. The formal worst-case calculation over all initial distributions also gives 13 steps for this benchmark at the same tolerance.

The spectral-gap reference scale is approximately 21 steps. This is reported as a conservative comparison scale and not as a universal finite-time theorem for the non-reversible benchmark. More importantly, the additional numerical experiments show that the diagnostic behavior changes substantially for slow-mixing, periodic, and nearly decomposable chains. The period-2 example confirms that graph-based cycle analysis is required to detect real negative unit eigenvalues, such as $\lambda = -1$.

The numerical study uses synthetic four-state examples. Reliability, deterioration, and risk modeling can be evaluated in future work using calibrated real-world transition matrices.

Future research should examine sparse and higher-dimensional matrices, uncertainty in estimated transition probabilities, non-homogeneous chains, and data-driven calibration. Future applied studies can also evaluate the framework using real reliability, deterioration, or risk-assessment datasets.

REFERENCES

[1] V. Knights, O. Petrovska and M. Prchkovska. "Mathematical analysis of stochastic transition matrices in markov models using monte carlo and MCMC simulation". *Journal of Natural Sciences and Mathematics of UT*, vol. 10, no. 19-20, pp. 100-110, 2025.

[2] L. Huzsvai, E. Kovács, G. Tuba, C. Juhász, D. Jug and J. Zsembeli. "Discrete-time markov chain method for predicting probability of crop yield variability". *Earth*, vol. 6, no. 4, p. 142, 2025.

[3] R. K. Gupta, D. Khan and S. Banerjee. "Markovian Brand switching model for long-term steady-state market shares: A study on toothpaste market". In: *Decision Making Under Uncertainty Via Optimization, Modelling, and Analysis*. Springer Nature, Singapore, pp. 567-581, 2025.

[4] D. S. S. Mohan, K. V. S. Chakresh, K. S. Reddy and Y. Sailaja. "A Real-Time System for Portfolio Optimization Using Markov Chain Steady-State Analysis". In: *2025 9th International Conference on Computational System and Information Technology for Sustainable Solutions (CSITSS)*. IEEE, pp. 1-6, 2025.

[5] A. Raju. "Applications of Markov chains in climate change modelling: A comprehensive review of advances, challenges, and future directions". *Ecological Modelling*, vol. 514, p. 111470, 2026.

[6] S. Han, Z. Grieser, S. Yoshikawa, T. Noro, T. Suda and K. Ho. "*Spare Strategy Analysis and Design for Large-Scale Satellite Constellation Using Markov Chain*". arXiv [preprint], 2025.

[7] W. Shao, T. Cai and X. Yan. "Advances on stability analysis for discrete-time Markovian jump impulsive stochastic systems with hybrid switching dynamics". *Neurocomputing*, vol. 652, p. 131128, 2025.

[8] J. Wang, M. Ye, Y. Kuang, R. Yang, W. Zhou, H. Li and F. Wu. "Long-term feature extraction via frequency prediction for efficient reinforcement learning". *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 47, pp. 3094-3110, 2025.

[9] A. Habib and M. W. Grenn. "The Effects of Digital Twins Development on System's Long-Term Performance, Potential Capabilities, and Possible Benefits". In: *International Conference on Enterprise Information Systems (ICEIS)*. pp. 300-307, 2025.

[10] A. Nazerian, M. Asllani, M. Tyloo and F. Sorrentino. "Open networks in discrete time: Passing vs blocking behavior". *Chaos: An Interdisciplinary. Journal of Nonlinear Science*, vol. 35, no. 12, p. 123113, 2025.

[11] S. Feng, L. Li and D. Yuan. "A markov chain monte carlo-based hybrid noise inference for continuous wavelet power spectra: With applications to solar and stellar oscillatory signals". *The Astrophysical Journal Supplement Series*, vol. 285, no. 1, p. 17, 2026.

[12] Z. T. Hornberger, D. M. King and S. H. Jacobson. "A framework for analyzing the periodically-observed time-homogeneous Poisson process". *Journal of the Operational Research Society*, vol. 76, no. 12, pp. 2532-2547, 2025.

[13] B. I. García-Maya, Y. Morales-Huerta and R. Salgado-García. "Disease spread model in structurally complex spaces: An open Markov chain approach". *Journal of Computational Biology*, vol. 32, no. 4, pp. 394-416, 2025.

[14] Q. Wang. "*Spectral Purification in Reversible Markov Chains: Hidden Parameters, Observable Equivalence, and Finite-Time Rigidity*". arXiv [preprint], 2026.

[15] A. Branquinho, A. Foulquié-Moreno and M. Mañas. "*Spectral Theory for Markov Chains with Transition Matrix Admitting a Stochastic Bidiagonal Factorization*". arXiv [preprint], 2026.

[16] A. Thiam. "*Statistical Limit Theorems for Axiom A Diffeomorphisms: Exponential Mixing, Central Limit Theorem, and Large Deviations*". arXiv [preprint], 2026.

[17] B. Nettasinghe, S. Chatterjee, R. Tipireddy and M. M. Halappanavar. "Extending Conformal Prediction to Hidden Markov Models with Exact Validity Via De Finetti's Theorem for Markov Chains". In: *International Conference on Machine Learning*. PMLR, pp. 25890-25903, 2023.

[18] H. Hennion and L. Hervé, (Eds). "*Limit Theorems for Markov Chains and Stochastic Properties of Dynamical Systems by Quasi-Compactness*". Springer, Berlin, Heidelberg, 2002.

[19] R. Benita, M. Elad and J. Keshet. "Spectral analysis of diffusion models with application to schedule design". *Advances in Neural Information Processing Systems*, vol. 38, pp. 2073-2127, 2026.

APPENDIX

Appendix A. MATLAB implementation of the spectral-computational framework.

The following listing provides the complete MATLAB implementation used to define the four matrices, validate their stochastic and graph structure, compute spectral diagnostics, generate Table 3, and generate Figures 1-6. The script is self-contained and uses the same numerical definitions reported in the manuscript.

```
clc;
clear;
close all;

% 1) Explicit transition matrices
P = [0.65 0.20 0.10 0.05;
0.15 0.55 0.20 0.10;
0.10 0.20 0.50 0.20;
0.05 0.15 0.25 0.55];

P_slow = [0.97 0.01 0.01 0.01;
0.01 0.97 0.01 0.01;
0.01 0.01 0.97 0.01;
0.01 0.01 0.01 0.97];

P_periodic = [0.00 0.50 0.00 0.50;
0.50 0.00 0.50 0.00;
0.00 0.50 0.00 0.50;
0.50 0.00 0.50 0.00];

P_near = [0.490 0.500 0.005 0.005;
0.500 0.490 0.005 0.005;
0.005 0.005 0.490 0.500;
0.005 0.005 0.500 0.490];

chains = {P,P_slow,P_periodic,P_near};
chainNames = {"Benchmark," "Slow mixing," "Periodic(period 2)," "Nearly decomposable"};

% 2) Initial condition and parameters
x_0 = [1 0 0 0];
nSteps = 500;
plotSteps = 60;
tol = 1e-12;
ε = 1e-3;

% 3) Analyze all chains
R = repmat(struct("π,"[], "eigvals," [], "ρ,"[], "gap,"[],"period," []...,
"irreducible," [], "worstTV," [], "mixingTime," []), 1, numel(chains));

for c = 1: numel(chains)
Pk = chains{c};
[n,m] = size(Pk);
if n ~= m
Error ("Transition matrix must be square.");
end
if any(Pk(:) < -tol)
Error ("Transition matrix contains negative entries.");
end
if any (abs(sum(Pk,2)-1) > 1e-10)
Error ("Each row of P must sum to 1.");
end

% Boolean reachability
A = Pk > tol;
reach = logical (eye(n)) | A;
for k = 1: n
for i = 1: n
if reach (i,k)
reach (i,:) = reach(i,:) | reach(k,:);
end
end
end
irreducibleFlag = all (reach(:));

% Graph-based period
dist = -ones (n,1);
dist(1) = 0;
queue = 1; head = 1;
while head <= numel(queue)
u = queue (head); head = head + 1;
nbr = find (A(u,:));
for q = 1: numel(nbr)
v = nbr(q);
if dist(v) < 0
dist(v) = dist(u) + 1;
queue (end+1) = v; %#ok<AGROW>
end
end
end
period = 0;
for i = 1: n
nbr = find (A(i,:));
for q = 1: numel(nbr)
j = nbr (q);
if dist (j) >= 0
```

```
period = gcd (period, abs (dist(i) + 1-dist (j)));
end
end
end
if period == 0, period = 1; end

% Spectral decomposition and stationary distribution
[V,D] = eig(Pk');
eigvals = diag(D);
[~,idx] = sort(abs(eigvals), "descend");
eigvals = eigvals(idx); V = V(:,idx);
[~, idx1] = min (abs(eigvals-1));
π_vec = real (V(:,idx1));
if sum (π_vec) < 0, π_vec = -π_vec; end
π_vec = max (π_vec/sum(π_vec),0);
π_vec = π_vec/sum(π_vec);
absEig = abs (eigvals);
ρ = absEig (2);
gap = 1-ρ;

% Worst-case TV over the probability simplex
Ppow = eye (n);
worstTV = zeros (nSteps+1,1);
for t = 0: nSteps
rowTV = 0.5*sum (abs(Ppow-π_vec'),2);
worstTV(t+1) = max(rowTV);
if t < nSteps, Ppow = Ppow*Pk; end
end
idxMix = find (worstTV <= ε,1, "first");
if isempty(idxMix), mixingTime = NaN; else, mixingTime = idxMix-1; end

R(c).π = π_vec; R(c).eigvals = eigvals; R(c).ρ = ρ;
R(c).gap = gap; R(c).period = period; R(c).irreducible = irreducibleFlag;
R(c).worstTV = worstTV; R(c).mixingTime = mixingTime;
end

% 4) Benchmark trajectory and Figures 1-4
x = x0; X = zeros(plotSteps+1,4); X(1,:) = x0;
for t = 1:plotSteps, X(t+1,:) = X(t,:)*P; end
π = R(1).π;
tvDist = zeros(plotSteps+1,1); l2Dist = zeros(plotSteps+1,1);
for t = 0:plotSteps
d = X(t+1,:)-π';
tvDist(t+1) = 0.5*sum(abs(d)); l2Dist(t+1) = norm(d,2);
end

figure ("Color," "w"); plot(0:plotSteps,X, "LineWidth,"1.8); hold on;
for i = 1:4, plot([0 plotSteps],[π(i) π(i)], "--," "LineWidth,"1.0); end
xlabel("Time step"); ylabel("Probability"); title("State Probability Evolution"); grid on;
exportgraphics(gcf, "Figure 1_State_Probability_Evolution.png," "Resolution,"600);

figure("Color," "w"); semilogy(0:plotSteps,tvDist, "o-," "LineWidth,"1.6); hold on;
semilogy(0:plotSteps,l2Dist, "s-," "LineWidth,"1.6); yline(ε, "--," "LineWidth,"1.0);
xlabel("Time step"); ylabel("Distance to stationarity"); title("Convergence Diagnostics"); grid on;
exportgraphics(gcf, "Figure 2_Convergence_Diagnostics.png," "Resolution,"600);

figure("Color," "w"); theta = linspace(0,2*π,500);
plot(cos(theta),sin(theta), "--," "LineWidth,"1.0); hold on;
plot(real(R(1).eigvals),imag(R(1).eigvals), "o," "MarkerSize,"7);
xlabel("Real part"); ylabel("Imaginary part"); title("Eigenvalue Spectrum"); axis equal; grid on;
exportgraphics(gcf, "Figure 3_Eigenvalue_Spectrum.png," "Resolution,"600);

meanReturnTimes = 1./π; figure("Color," "w");
yyaxis left; bar(1:4,π); ylabel("Stationary probability");
yyaxis right; plot(1:4,meanReturnTimes, "o-," "LineWidth,"1.8); ylabel("Mean return time");
xlabel("State index"); title("Stationary Probabilities and Mean Return Times"); grid on;
exportgraphics(gcf, "Figure 4_Stationary_Return_Times.png," "Resolution,"600);

% 5) Table 3 and Figures 5 and 6
ρVals = zeros (numel(chains),1); gapVals = zeros(numel(chains),1);
periodVals = zeros (numel(chains),1); mixVals = strings(numel(chains),1);
for c = 1: numel (chains)
ρVals(c) = R(c).ρ; gapVals(c) = R(c).gap; periodVals(c) = R(c).period;
if isnan(R(c).mixingTime), mixVals(c) = "No finite mixing time"; else, mixVals(c) = string(R(c).mixingTime); end
end
```

```
Table 3 = table (string (chainNames”), ρVals, gapVals, periodVals, mixVals.,
“Variable Names,” {“Chain,” “ρ,” “γ,” “Period,” “WorstCaseMixingTime”});
disp (Table 3);

figure (“Color,” “w”); hold on;
for c = 1: numel (chains)
semilogy (0: nSteps, max(R(c).worstTV, realmin), “LineWidth,”1.7);
end
yline (ε, “--,” “LineWidth,”1.0);
xlabel (“Time step”); ylabel(“Worst-case total variation distance”);
title (“Worst-Case Total-Variation Convergence”); grid on;
legend (chainNames, “Location,” “best”);
exportgraphics (gcf, “Figure 5_Worst_Case_TV.png,” “Resolution,”600);

figure (“Color,” “w”); bar (ρVals);
set (gca, “XTick,”1: numel (chains), “XTickLabel,” chainNames);
ylabel (“Subdominant spectral radius ρ”); title (“Subdominant Spectral Radius across Structural Regimes”); grid on;
exportgraphics (gcf, “Figure 6_Subdominant_Rho.png,” “Resolution,” 600);
```