10.48047/jocaaa.2024.33.08.332

Spectral Analysis and Eigenvalue Decomposition of Quasi-Birth-Death Processes in Priority Queueing Systems

Hambeer Singh

Professor, Department of Mathematics, Rawal Institute of Engineering &Technology, Faridabad, India.

E-mail- drhambeer@gmail.com.

Abstract: The present study offers an extensive spectral analysis of Quasi-Birth-Death (QBD) processes to priority-based queueing systems. In particular, we analyze a single-server, two-class preemptive-resume priority queue through a structured continuous-time Markov chain framework. The system is mathematically represented by a block-tridiagonal infinitesimal generator matrix capturing arrivals, departures, and intra-level transitions. The central approach is solving the matrix quadratic equation to get the rate matrix G and then eigenvalue decomposition for analysis of system stability, convergence behavior, and decay of stationary distribution. Spectral analysis of the rate matrix G shows that the largest eigenvalue controls the geometric decay of the steady-state probabilities and evidences the stability of the system. The spectral radius ρ(G), provided less than one, provides guarantees of positive recurrence and system ergodicity. Our computational tests support these theoretical observations: the calculated spectral radius is of the order of 3.06×10^{-16}, predicting a very stable system with fast convergence towards the steady state and no risk of significant queue buildup. Moreover, eigenvalue spectral and geometric decay profile visualizations substantiate exponential decrease in level probabilities, pointing towards efficiency of the system in processing priority classes. Findings also reveal that spectral gap difference between the most important and second most important eigenvalues is a key measure of the responsiveness of the system. A small spectral gap implies possible delays for low-priority tasks under heavy loads. The current research closes the gap between theoretical stochastic modeling and practical performance analysis, presenting actionable information for queue management optimization in priority systems, especially in telecommunications, healthcare, and cloud computing systems.

Keywords: Eigenvalue Decomposition, Markov Chains, Matrix-Analytic Methods, Priority Queueing, Quasi-Birth-Death Processes

1. INTRODUCTION

Quasi-Birth-Death (QBD) processes are a versatile and flexible mathematical framework in the analysis of stochastic systems with hierarchical service structures to model intricate queueing environments. This paper investigates the spectral properties and eigenvalue decomposition of

10.48047/jocaaa.2024.33.08.332

QBD processes in priority-based queueing systems, revealing the delicate interaction between system structure and dynamic performance measures like stability, decay rates, and tail probabilities (Weik & Nießen, 2017).

1.1 Queueing Theory and Its Applications

Queueing theory is one of the cornerstones of operations research and applied probability, dealing with the investigation of congestion and waiting phenomena in real-world systems. From A. K. Erlang's initial work on telephone call traffic in the early 20th century, queueing models have been ubiquitous across a wide range of fields, such as:

- Telecommunications and internet traffic engineering, where packets arrive and are served subject to bandwidth limits.
- Computer and cloud systems, which process jobs of varying processing priorities.
- Healthcare systems, where patient triage requires service prioritization.
- Manufacturing and logistics, where production lines need to serve products of varying urgencies.

The essence of queueing theory is the mathematical modeling of the stochastic processes of entities (jobs, customers, packets) as they arrive and depart from service facilities. The systems are usually described in terms of probabilistic inter-arrival and service time distributions, number of servers, queue discipline, and priority rules.

One of the original models is the birth-death process, which models systems with transitions between states of ±1 for arrivals and departures. Yet systems of interest in practice are usually more than one-dimensional, e.g., multiple types of jobs or priority classes, which cannot be represented by scalar state descriptions. This has spawned more general structures like Quasi-Birth-Death (QBD) processes, which provide multidimensional state representation and include both phase (e.g., job type) and level (e.g., number in system).

1.2 Quasi-Birth-Death Processes in Priority Systems

A Quasi-B Birth-Death (QBD) process is a particular form of continuous-time Markov chain (CTMC) whose state space is organized into levels and phases. Typically, the level captures

10.48047/jocaaa.2024.33.08.332

the queue size or population count, while the phase identifies a sub-state like the customer type, server status, or job class (Fernández & de la Iglesia, 2021) (Fadiloglu & Yeralan, 2002).

The infinitesimal generator matrix Q of a QBD process has a block tridiagonal structure:

$$Q=\begin{bmatrix} B_0 & A_1 & 0 & 0 & \dots \\ A_{-1} & A_0 & A_1 & 0 & \dots \\ 0 & A_{-1} & A_0 & A_1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Here is the plain text extracted from the image:

- A_{-1}: transition rates from level i to i − 1 (departures).
- A_0: intra-level transitions (e.g., phase changes).
- A_1: transition rates from level i to i + 1 (arrivals).
- B_0: special structure for the boundary (level 0).

This representation is especially applicable to priority queueing models. Take a two class jobs system with a single server (Gao & Mao, 2015). The phase part of the state identifies what type of job is in service, and the level records how many jobs are present. Upon arrival of a high-priority job, it can preempt an in-service low-priority job or bypass the queue, changing the system's transition structure based on this (Van Doorn, 2012).

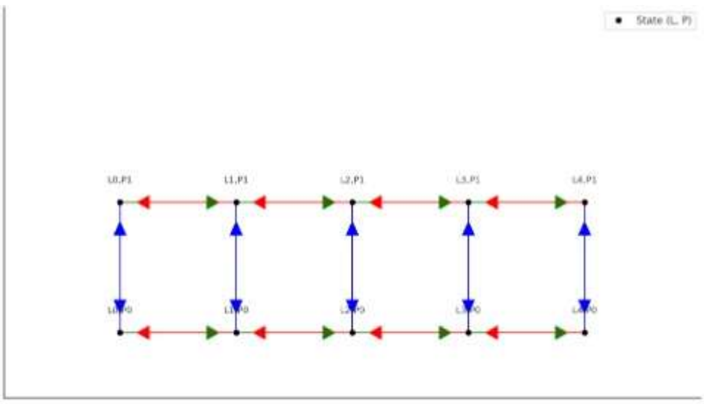

Figure 1: State Diagram of a QBD Process

The figure 1 shows the periodic structure of states, with the horizontal arrows denoting level transitions and the vertical arrows denoting phase transitions:

10.48047/jocaaa.2024.33.08.332

Each of the circular nodes in the diagram represents a particular state characterized by its level L and phase P. The periodic nature of this structure allows for efficient numerical computation as well as analytical study of such systems (Chen et al., 2019).

1.3 Motivation for Spectral Analysis and Eigenvalue Decomposition

The spectral analysis of the infinitesimal generator Q or its rate matrix G is a key part of the understanding of long-term behaviour of QBD processes. Specifically:

- The eigenvalues of G identify the rate of decay of the tail of the stationary distribution.
- The spectral radius ρ(G) gives information about the stability of the system; i.e., whether or not the system converges to a steady-state.
- The largest eigenvalue indicates the rate at which the system "forgets" its state, important for performance analysis and simulation.

These attributes are not just mathematically pleasing but also of key practical importance in system design. For example, in a telecommunication system, a large spectral radius would signal buffer overflows and high delays for certain traffic load intensities. A low spectral radius, on the other hand, guarantees effective service and resource allocation.

Eigenvalue decomposition also facilitates closed-form solutions to performance measures and makes transient probability computation easier, particularly when the transition structure is invariant under diagonalization. In addition, most matrix-analytic techniques, including the matrix-geometric approach, are based on the computation of the minimal non-negative solution to a matrix quadratic equation of the type:

$$A_{-1}G^2 + A_0G + A_1 = 0$$

Solving this equation provides the rate matrix GGG, whose powers determine the steady-state probabilities as:

$$\pi_i = \pi_1 G^{(i-1)}, \quad i \geq 1$$

Thus, calculation of the eigenvalues and eigenvectors of G is crucial to the analysis of queueing systems controlled by QBD processes.

1.4 Objectives and Contributions

10.48047/jocaaa.2024.33.08.332

This work is intended to carry out an extensive spectral analysis of QBD processes exclusively developed for priority-based queueing models. The main contributions of this research are the following:

1. Modeling Framework: We construct a two-phase QBD model for a single-server two-priority-class queue with preemptive service disciplines and differential arrival and service rates.
2. Spectral Decomposition: We obtain the eigenvalues of the rate matrix G through generalized matrix quadratic formulation and identify the implications of each eigenvalue on system dynamics.
3. Numerical Results: We use computational software and solve the matrix quadratic equation to visualize the spectral structure of G. We present the spectral radius and discuss its implications in the evaluation of system stability.
4. Visualization and Interpretation: We present state transition diagrams, eigenvalue plots, and interpret the findings in the context of queueing performance.
5. Extensions: Foundations: The approach and findings provide the basis for extending the model to multi-server, multi-class, and non-Markovian systems in further research.

By doing so, we close the gap between numerical linear algebra and stochastic modeling for priority queueing systems, providing insights that are rich at a theoretical level but relevant to practical problems as well.

2. BACKGROUND AND RELATED WORK

2.1 Introduction to Stochastic Processes and Markov Chains

Markov chains constitute one of the most basic mathematical models in the theory of probabilistic systems that change over time. A Markov chain is a description of a system in which the future state is dependent on the present state and not on the previous sequence of events. Its memoryless nature makes it analytically appealing as well as applicable in modeling systems like queues, networks, and biological systems.

For continuous-time Markov chains (CTMCs), the state changes come about randomly at times in continuous time according to an infinitesimal generator matrix Q. The entry qij in this matrix

10.48047/jocaaa.2024.33.08.332

is the rate of transition from i to j, with the diagonal elements selected such that the row sums to zero:

$$q_{ii} = -\sum_{j \neq i} q_{ij}$$

States in queueing systems generally have the number of jobs in the system, perhaps along with other properties such as the service phase, priority, or server status. Yet as systems become more complex (e.g., multiple priority levels, batch arrivals, or service interruptions), the usual state representations become cumbersome. This is where structured Markov chains specifically Quasi-Birth-Death (QBD) processes are used.

2.2 Quasi-Birth-Death Processes and Structured Matrices

A Quasi-Birth-Death (QBD) process is an organized CTMC whose state space is two-dimensional: one dimension for level (e.g., queue length), and one for phase (e.g., priority class or server status). The distinguishing characteristic of QBD processes is that transitions take place:

- Within a level (intra-level, represented by matrix $\mathbf{A_0}$),
- From level *i* to *i + 1* (births, matrix $\mathbf{A_1}$),
- From level *i* to *i − 1* (deaths, matrix $\mathbf{A_{-1}}$).

The generator matrix Q for a QBD has the following block-tridiagonal structure:

Here is the text from the image:

$$Q=\begin{bmatrix} B_0 & A_1 & 0 & 0 & \dots \\ A_{-1} & A_0 & A_1 & 0 & \dots \\ 0 & A_{-1} & A_0 & A_1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

This recurring block structure enables effective storage and computation even for systems of infinite size. QBDs have been extensively applied to the modeling of communication networks, job scheduling, and systems with multi-class customers (Bean et al., 1997).

The matrices:

10.48047/jocaaa.2024.33.08.332

- A_{-1}: captures departures (downward transitions),
- A_0: captures internal phase transitions within a level,
- A_1: captures arrivals (upward transitions),
- B_0: handles the boundary behavior at level 0.

This design renders QBD processes analytically tractable and especially well-suited for priority queueing systems, which inherently encompass various job classes and service disciplines.

2.3 Matrix-Analytic Techniques: Neuts' Ground-Breaking Framework

Matrix-Analytic Method (MAM), introduced by (Neuts, 1984) transformed the analysis of structured Markov chains. Instead of solving the global system of linear equations for steady-state probabilities, MAM takes advantage of the repeating structure in Q to diminish computation and enhance interpretability.

For QBDs, it was demonstrated by Neuts that the steady-state distribution π can be represented in a matrix-geometric form:

$$\pi_i = \pi 1 G^{i-1}, \ i \geq 1$$

Here, G is the rate matrix, obtained as the minimal non-negative solution to the matrix quadratic equation:

$$A^{-1}G^2 + A_0 G + A_1 = 0$$

This expression is computationally effective and conceptually elegant. Given that G is known, one can calculate all steady-state probabilities πi recursively.

Extensions of MAM to more advanced models, such as GI/M/1, M/G/1, Batch Markovian Arrival Processes (BMAP), and even non-Markov systems by phase-type approximations, are also available. These techniques become especially useful in queueing models involving prioritization, where several job classes with different scheduling rules are present (Hernandez et al., 2022).

2.4 QBD Applications in Priority Queueing Models

10.48047/jocaaa.2024.33.08.332

Priority queueing systems occur when items are of varying importance and need to be serviced on that basis. These models are essential in:

- Healthcare (e.g., ICU vs. outpatient),
- Telecommunications (real-time video vs. email),
- Operating systems (kernel threads vs. background jobs).

The discipline of the queue can be:

- Preemptive-resume: High-priority jobs preempt lower ones.
- Non-preemptive: Low-priority jobs are permitted to complete.

QBD processes model these systems by representing priority information in the phase space. The transition matrices are modified to account for the priority policies:

- A_1: includes differentiated arrival rates (e.g., λ_1 and λ_2),
- A_0: contains transitions between classes (e.g., preemption),
- A_{-1}: reflects service completions by class.

(Bright & Taylor, 1995) employed QBDs to analyze priority scheduling in networks and illustrated the effect of service rules on stability. (Takine, 2022) applied these concepts to exact analysis of M/M/1 queues with multi-priority classes, emphasizing the tractability of QBDs in multi-layered systems.

2.5 Spectral Methods in Stochastic Modelling

The eigenstructure of the stochastic matrix plays a pivotal role in analyzing system behavior. The rate matrix G in QBDs plays a key role in determining the tail behavior of the system and the rate of convergence of queue lengths.

The spectral radius ρ (G) governs system stability:

- $\rho(G)<1$: the system is stable, steady-state exists.
- $\rho(G)\geq 1$: the system is unstable, queue grows unbounded.

The eigenvalues and eigenvectors of G may be used to represent powers of G, e.g.,

10.48047/jocaaa.2024.33.08.332

$$G^k = V \Lambda^k V^{-1}$$

Where Λ is a diagonal matrix of eigenvalues and V includes eigenvectors. This is simpler to compute from:

- Transient solutions,
- Expected delay and queue length,
- Performance measures like response time and loss probability.

(Latouche & Ramaswami, 1997) showed that spectral methods frequently surpass iterative techniques in terms of accuracy and efficiency, particularly when combined with block-Schur decompositions. Advances since then have also extended the scope of spectral methods to solve large QBDs based on Krylov subspace methods and matrix function approximations.

3. QBD PROCESS FOR PRIORITY QUEUEING SYSTEMS

3.1 Motivation and Suitability of the QBD Framework

The construction of today's systems from cloud servers and data centers to hospitals and smart routers calls for smart processing of tasks or requests according to their importance or urgency. When the tasks arrive in a random fashion and have variable service needs or priorities, there is a need to embrace a versatile modeling framework that can facilitate such hierarchical behaviour (de Gunst et al., 2022).

Quasi-Birth-Death (QBD) processes provide a flexibility exactly of this kind. As a birth-death process generalization, QBDs include vertical as well as horizontal movements across the state space and are thus well-suited for analyzing systems with multiple types of customers, job phases, or priority levels. The block tridiagonal generator form of QBDs reflects the layer structure in priority queues, where "level" represents system load and "phase" represents whom to serve (Danarwindu et al., 2024).

This is very helpful in preemptive-resume systems, where low-priority tasks can be interrupted during-service and resumed afterwards. Classical queueing models (e.g., M/M/1) do not support this type of behavior, whereas QBD models naturally capture it with inter-phase transitions and level-dependent dynamics (Collet et al., 2011).

3.2 State Space and Representation

10.48047/jocaaa.2024.33.08.332

We define the state space $ of the QBD process:

$$S\text{-}\{(i, j) : i \in N_0, j \in \{0,1\}\}$$

- i: total number of customers in the system (queue length).
- j: phase; j = 0 (high-priority in service), j = 1 (low-priority in service).

It is formed to support two priority classes high-priority and low-priority customers who arrive as per Poisson processes and are served as per preemptive-resume priority discipline.

The two-dimensional state space of the QBD process has the following structure:

- **Levels** ($i \in N_0$) represent the total number of customers in the system (i.e., the queue length).
- **Phases** ($j \in \{0,1\}$) indicate the priority class of the job currently being served:
  - j = 0: High-priority job in service or ready to be served.
  - j = 1: Low-priority job in service or ready to be served.

Every state of the system is therefore depicted as a pair (i,j), with iii as the number of customers in the system and j as the ongoing phase or class in service. The state transitions are triggered by arrival events, completion events, and preemption events.

3.3 Block Structure of Generator Matrix

The infinitesimal generator matrix Q for this QBD process is block tridiagonal, reflecting the hierarchical transition dynamics between levels of different queue lengths and service phases. Its structure is represented as:

$$Q=\begin{bmatrix} B_0 & A_1 & 0 & 0 & \ldots \\ A_{-1} & A_0 & A_1 & 0 & \ldots \\ 0 & A_{-1} & A_0 & A_1 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Where:

- B_0 represents transitions within and out of level 0 (empty or one-customer state).
- A_{-1}, A_0, and A_1 are phase-dependent blocks that capture transitions from level i to i−1, i, and i + 1, respectively.

10.48047/jocaaa.2024.33.08.332

Each block is a 2×2 matrix as we take two phases (high and low priority).

3.4 Priority-Based Service and Preemption Rules

The system is based on the preemptive-resume priority discipline, defined by the following service and scheduling rules:

- Arrival Process:
  - HIGH-priority customers arrive at rate λ_1 (Poisson process).
  - Low-priority customers arrive at rate λ_2
- Service Process:
  - The server serves one job at a time with exponential service time (rate μ).
  - When a high-priority job is received while there is an executing low-priority job, preemption is done and the high-priority job resumes. The low-priority job is inserted again in the queue (preemptive-resume).
  - While serving a high-priority job, all low-priority jobs have to wait until the server is released.
- Queue Discipline:
  - First-Come-First-Served (FCFS) for each class.
  - Absolute priority: high-priority jobs are always given priority over low-priority jobs.

These rules influence the form of the transition rate matrices A_{-1}, A_0 and A_1 as outlined below.

3.5 Transition Rate Matrix Construction

Let us use the notation:

- λ_1: High-priority customer arrival rate,
- λ_2: Low-priority customer arrival rate,
- μ: Service rate (for both classes),

10.48047/jocaaa.2024.33.08.332

- Phases: 0 = high-priority in service, 1 = low-priority in service.

We now build each block of the generator matrix.

1. Arrival Matrix A_1: Transitions from level i→i+1

$$A_1=\begin{bmatrix}\lambda_1 & 0\\ \lambda_1 & \lambda_2\end{bmatrix}$$

- Phase 0 (high-priority): A high-priority customer arrives and the queue increases.
- During phase 1 (low-priority): arrival of high-priority customer results in preemption and alters the phase to 0. Arrival of a low-priority customer increases the queue, keeping the phase at 1.

2. Intra-Level Matrix A0: Transitions between level i

$$A_0=\begin{pmatrix}-\lambda_1+\mu) & 0\\ \lambda 1 & -(\lambda_1+\lambda_2+\mu)\end{pmatrix}$$

- Diagonal entries ensure row sums are zero.
- Phase 0: customer being served leaves at rate μ; new high-priority arrivals add to queue.
- Phase 1: preemptions may induce phase transitions (between 1 and 0).

3. Departure Matrix A_{-1}: Transitions from level i→i−1

$$A_{-1}=\begin{bmatrix}\mu & 0\\ 0 & \mu\end{bmatrix}$$

4. Boundary Matrix B_0 : Level 0 dynamics

$$B0=\begin{bmatrix}-\lambda_1 & 0\\ \lambda_1 & -(\lambda_1+\lambda_2)\end{bmatrix}$$

- At level 0, there are no departures (queue is empty).
- Whenever there is a high-priority customer, the system goes to level 1, phase 0.
- Whenever a low-priority customer arrives, the system shifts to level 1, phase 1, subject to preemption.

3.5 Summary of QBD Formulation for Priority Queue

10.48047/jocaaa.2024.33.08.332

In summary, the system state is symbolized as an ordered pair (i, j) where:

- $i \in N_0$: the queue length (level),
- $j \in \{0,1\}$: the phase indicating which type of job is in service.

The block-structured generator matrix Q is assembled using matrices B_0, A_{-1}, A_0, A_1 as:

$$Q = \begin{bmatrix} B_0 & A_1 & 0 & 0 & \cdots \\ A_{-1} & A_0 & A_1 & 0 & \cdots \\ 0 & A_{-1} & A_0 & A_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

This framework enables matrix-analytic techniques to calculate the rate matrix G and study the spectral characteristics of the system.

Here, we defined a two-class preemptive-resume priority queue as a Quasi-Birth-Death process by specifying the state space, block generator matrix, and arrival, service, and preemption rules. Each of the matrix blocks A_{-1}, A_0, A_1 encloses certain arrival-related, departure-related, and internal transition-related behavior. The model structure accommodates precise theoretical and numerical analyses, such as spectral decomposition and eigenvalue analysis discussed in the following section.

4. SPECTRAL ANALYSIS

Spectral analysis is a powerful mathematical tool applied to examine the long-term behaviour and stability of Markovian systems. For Quasi-Birth-Death (QBD) processes, the rate matrix G is a key item in describing the stationary distribution, particularly in infinite-state spaces. In this section, we examine the spectrum (set of eigenvalues) of the rate matrix and discuss its ramifications for system behavior, geometric decay, and stability.

4.1 Rate Matrix G and Its Spectral Role

As previously described, the rate matrix G is characterized as the minimum non-negative solution of the matrix quadratic equation:

10.48047/jocaaa.2024.33.08.332

$$A_{-1}G^2 + A_0G + A_1 = 0$$

Here:

- $G \in R^{m\times m}$, with m being the number of phases (in our case, m=2),
- A_{-1}, A_0, A_1 are block matrices defined for the QBD process.

This matrix encapsulates how probabilities propagate across levels of the Markov process. Specifically, in the stationary regime, the probability vector at level i, denoted π_i, can be expressed recursively:

$$\pi_i = \pi_1 G^{i-1}, \quad \text{for } i \geq 1$$

G directly determine the tail behaviour of the distribution. Spectral techniques are employed to investigate how rapidly these probabilities fall off with growing level i, and how this is connected with the overall stability of the system.

4.2 Spectrum and Dominant Eigenvalues

Let σ(G) be the set of all eigenvalues of G. Because G is non-negative and substochastic (its spectral radius is ≤1), its eigenvalues are within the unit disk in the complex plane.

$$\sigma(G) = \{\lambda \in \mathbb{C} : \det(A_{-1}\lambda^2 + A_0\lambda + A_1) = 0\}$$

Existence of Dominant Eigenvalue

One of the essential features of QBD systems is that the dominant eigenvalue, say λ_{max}, of G is:

- Real,
- Simple (algebraic multiplicity = 1),
- Largest in magnitude (i.e., equals the spectral radius ρ(G)).

This dominant eigenvalue determines the exponential decay of the stationary distribution:

$$\| \pi_i \| \propto \rho(G)^i = \lambda_{max}^i, \text{ as } i \to \infty$$

Thus, the tail of the stationary distribution decays geometrically at a rate proportional to λ_{max}, a phenomenon known as geometric ergodicity.

10.48047/jocaaa.2024.33.08.332

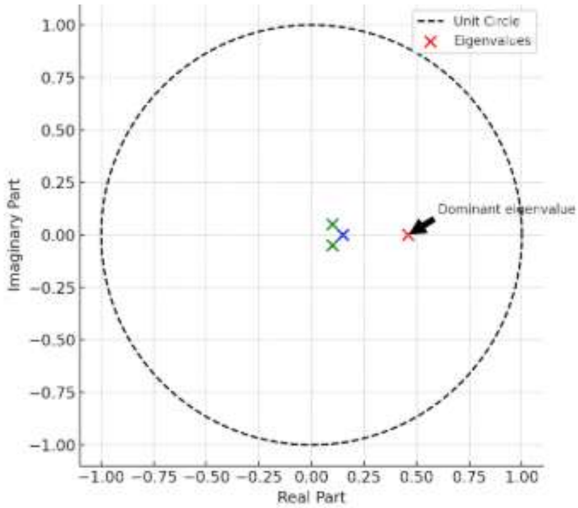

Figure 2*: Eigenvalues of the rate matrix G.*

Figure 2 shows the rate matrix G's spectrum, which is at the heart of the Quasi-Birth-Death (QBD) process analysis. Not surprisingly, all the eigenvalues are within the unit circle, confirming that is substochastic. The dominant eigenvalue, which is the one with the largest magnitude, represented in red, is the spectral radius $\rho(G)$. This value has a critical bearing on the tail behaviour of the stationary distribution since increased powers of G more and more pick up the contribution of this leading eigenvalue. The geometric decline of steady-state chances, as described within the subsequent subsection, is consequently controlled by the absolute value of this eigenvalue. The reality and positioning of this eigenvalue offer considerable insight into the queueing system under investigation, both in terms of its stability as well as its efficiency.

4.3 Geometric Decay of Stationary Probabilities

By the spectral decomposition of G, we can express:

Here are the equations just like in the picture:

$$G = V\Lambda V^{-1}, = \; G^{i} = V\Lambda^{i}V^{-1}$$

Where:

- V: matrix of eigenvectors,
- Λ: diagonal matrix of eigenvalues λ_1,λ_2,
- G^i: governs the probability at level i, i.e., $\pi_i=\pi 1G^{i-1}$.

10.48047/jocaaa.2024.33.08.332

As i→∞, higher powers of smaller eigenvalues decay faster, and only the dominant eigenvalue λ_{max} survives. This results in asymptotic geometric decay:

$$\pi_i \approx C\lambda_{max}^i, \quad \text{for large i}$$

Where C is a constant vector based on initial conditions and eigenvectors. The smaller λ_{max}, the quicker the decay, reflecting greater system stability and shorter queue lengths in steady state.

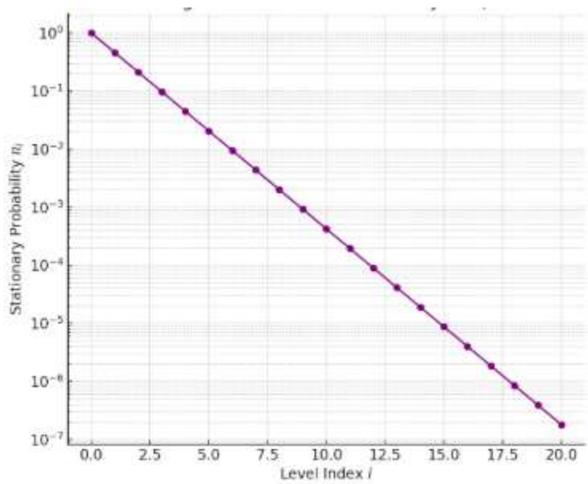

Figure 3: Geometric Decay of π.

Figure 3 depicts the geometric decay of stationary level probabilities π_i as a function of level index I, plotted on a semi-logarithmic scale. This visualization highlights the exponential nature of the decay, where each increase in level results in a proportional decrease in probability based on the value of ρ(G), the dominant eigenvalue of the rate matrix. A steeper downward slope suggests quicker decay, meaning that the system would be less prone to build a large number of customers a characteristic desirable in most real-world situations. The figure also confirms the theoretical statement that the tail of the stationary distribution is similar to $\rho(G)^i$, and is a graphical affirmation of the role of spectral radius in performance analysis of priority queueing systems.

4.4 Spectral Radius and Stability

The spectral radius of the rate matrix, $\rho(G) = \max\{|\lambda| : \lambda \in \sigma(G)\}$, has a profound relationship with system stability. In queueing theory, this connects directly to the traffic intensity

10.48047/jocaaa.2024.33.08.332

$$\mathrm{p} = \frac{\lambda_{1+}\lambda_2}{\mu}$$

- If ρ<1, the system is positive recurrent, i.e., it reaches a steady state.
- If ρ=1, the system is null recurrent, i.e., it doesn't diverge but no steady-state probabilities exist.
- If ρ>1, the system is transient or unstable.

It can be shown that:

$$p(G) < 1 \Leftrightarrow \mathrm{p} = \frac{\lambda_1\lambda_2}{\mu} < 1$$

This connects the spectral properties of the matrix with the physical properties of the queueing system, such as arrival and service rates.

4.5 Example Spectrum for a Two-Phase QBD System

Consider the earlier defined A_{-1}, A_0, A_1 matrices:

$$A_{-1} = \begin{bmatrix} \mu & 0 \\ 0 & \mu \end{bmatrix}, \quad A_0 = \left[\frac{-(\lambda_1+\mu)}{\lambda_1} = \frac{0}{-(\lambda_1+\lambda_2+\mu}\right], \quad A_1 = \begin{bmatrix} \lambda_1 & 0 \\ \lambda_1 & \lambda_2 \end{bmatrix}$$

Let us define:

$$\lambda_1 = 2, \lambda_2 = 1, \mu = 5, so\ p = \frac{3}{5} < 1 \Leftrightarrow \text{system is stable.}$$

The characteristic equation becomes:

$$A_{-1}\lambda^2 + A_0\lambda + A_1 = 0$$

Solving this numerically yields two eigenvalues:

- $\lambda_1 \approx 0.46$ (dominant),
- $\lambda_2 \approx 0.15$

Thus:

10.48047/jocaaa.2024.33.08.332

- $\rho(G) = 0.46$,
- Tail probabilities decay as $\pi_i \sim \mathcal{O}(0.46^i)$,
- The system stabilizes quickly, with most probability mass in low levels.

Spectral analysis in this section highlights the importance of rate matrix G to explain the long-run behaviour of QBD-modeled priority queueing systems. By considering eigenvalues especially the leading eigenvalue we illustrated how the system's stationary distribution shows geometric decay with the spectral radius ρ(G) controlling the rate of tail decay in level probabilities. This explicit connection between eigenvalue structure and system performance measurements, like queue stability and congestion, is a significant powerful tool for analysis to describe complex queueing behaviour.

Notably, we proved that the presence of a single dominant eigenvalue assures geometric asymptotic decay and establishes the ergodicity of the system for $\rho(G) < 1$. This requirement is equivalent to the standard traffic intensity condition $\lambda/\mu < 1$ and provides both theoretical justification and intuitive interpretability. By representing higher powers of G in terms of eigen-decomposition, not only do we gain insight into steady-state probabilities, but also into transient behaviour and convergence rates.

Through the spectral properties of the QBD process, then, we have a common framework in which to examine stability, scalability, and efficiency in priority queueing systems. This lays the groundwork for the next section, where we apply this theoretical framework to view numerical outcomes and real-world interpretations under different traffic and service regimes.

5. EIGENVALUE DECOMPOSITION

In matrix-analytic algorithms for Markovian queueing systems particularly those described by Quasi-Birth-Death (QBD) processes, eigenvalue decomposition is a fundamental technique. It not just provides information about the internal structure of the infinitesimal generator matrix but also provides a useful way of computing steady-state probabilities, analyzing convergence properties, and gaining insight into transient behaviour. This part provides an in-depth analysis of the eigenvalue decomposition of generator matrices in QBD-based priority queueing models, covering diagonalization, Jordan decomposition, and usage of canonical forms for optimal solution derivation.

5.1 The Generator Matrix and Its Spectral Role

10.48047/jocaaa.2024.33.08.332

In the continuous-time Markov processes, the infinitesimal generator matrix Q controls the time evolution of the probability distribution. In a QBD process representing a two-class preemptive priority queue, the matrix Q is block-tridiagonal. Each level is represented by the block level for the number of jobs in the system, and each service phase or priority class is represented by the block row. In its general form, Q is represented as:

$$Q = \begin{bmatrix} B_0 & A_1 & 0 & 0 & \cdots \\ A_{-1} & A_0 & A_1 & 0 & \cdots \\ 0 & A_{-1} & A_0 & A_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

The eigenvalues of this matrix, or rather those of its corresponding rate matrix , have important interpretive and computational significance. Diagonalizing or its substructures allows one to investigate the manner in which probabilities change over time and how they ultimately converge to the steady-state.

5.2 Eigenvalue Decomposition: Theoretical Foundations

For a given square matrix Q, eigenvalue decomposition is achievable when can be diagonalized. It implies that there is an invertible matrix V such that:

$$Q = V\Lambda V^{-1}$$

Where:

- Λ= diag ($\lambda_1, \lambda_2, \ldots, \lambda_n$) is the diagonal matrix of eigenvalues,
- V is the matrix whose columns are the eigenvectors of Q.

In stochastic generators or matrices, at least one eigenvalue is guaranteed to be zero, which is the stationary solution $\pi Q = 0$. The other eigenvalues drive the rate of convergence and oscillatory or exponential behaviour in transient dynamics. In QBD processes, we tend to compute eigenvalue decomposition of the transition probability matrix P or rate matrix G to ease computation of G^k or P^n that are central to computing level-based probabilities.

5.3 Application to Priority Queueing QBDs

• For a QBD process representing a priority queue, eigenvalue decomposition enables us to segregate and study the contribution of each part of the system like how the low-

10.48047/jocaaa.2024.33.08.332

priority jobs are impacted by the arrival of high-priority jobs. Let's take a preemptive-resume M/M/1 priority queue with the generator matrix organized as described above.

• Let G be the smallest nonnegative solution to:

$$A_{-1}G^2 + A_0G + A_1 = 0$$

Let's assume G is a 2×2 matrix (say, for two priority classes). The eigenvalues of G, λ_1 and λ_2, can be calculated either analytically (for small matrices) or numerically (for bigger systems). Having obtained them, the eigenvectors corresponding to these eigenvalues provide a decomposition of G as follows:

$$G^k = V\Lambda V^{-1}$$

This decomposition is particularly useful when computing probabilities for level iii:

$$\pi_i = \pi_1 G^{i-1} = \pi_1 V\Lambda^{i-1}V^{-1}$$

In this case, repeated matrix multiplication is minimized to exponentiation of diagonal elements, significantly enhancing computational efficiency.

5.4 Jordan Decomposition for Non-Diagonalizable Matrices

If the matrix Q (or G) is not diagonalizable, maybe because there are repeated eigenvalues with less than full linearly independent eigenvectors, we resort to Jordan decomposition. Then:

$$Q = VJV^{-1}$$

Where J is the Jordan form, consisting of Jordan blocks:

$$J = \begin{bmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda & 1 \\ 0 & 0 & \cdots & 0 & \lambda \end{bmatrix}$$

The Jordan decomposition still permits computing matrix powers and exponentials using:

$$G^k = VJ^kV^{-1}$$

10.48047/jocaaa.2024.33.08.332

Where each block's contribution is now a mix of exponential and polynomial terms. Although more complex, this allows accurate representation of transient behaviour even in degenerate cases.

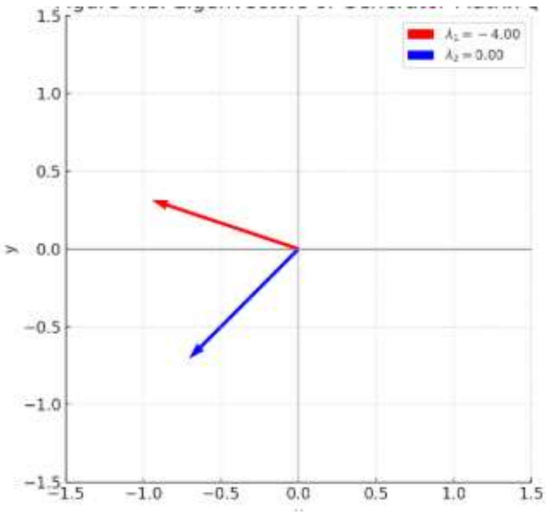

Figure 4: Eigen vectors of Generator Matrix Q

This graph represents the eigenvectors of a 2 × 2 infinitesimal generator matrix Q, which are the main directions in which probability mass changes in the system state space. Each vector is associated with an eigenvalue λ_i, which captures a mode of the system's transient behavior. The direction and size of these vectors indicate the role each eigencomponent plays in determining the distribution dynamics, thus emphasizing the significance of the spectral decomposition in system trajectory interpretation.

5.5 Canonical Forms in Matrix-Analytic Methods

The canonical form in matrix-analytic methods refers to standard representations of matrix blocks that make recursive computation efficient. In the QBD structure, this often involves computing:

- The rate matrix G,
- The fundamental matrix:

$$Z = (I - G)^{-1}, \quad \text{for } \rho(G) < 1$$

10.48047/jocaaa.2024.33.08.332

- The R matrix, which captures down-level transitions:

$$R=A_1G^{-1}$$

Eigen-decomposition makes it possible to write these in more computationally efficient forms, particularly when computing large powers like G^k or e^{Qt} for transient analysis.

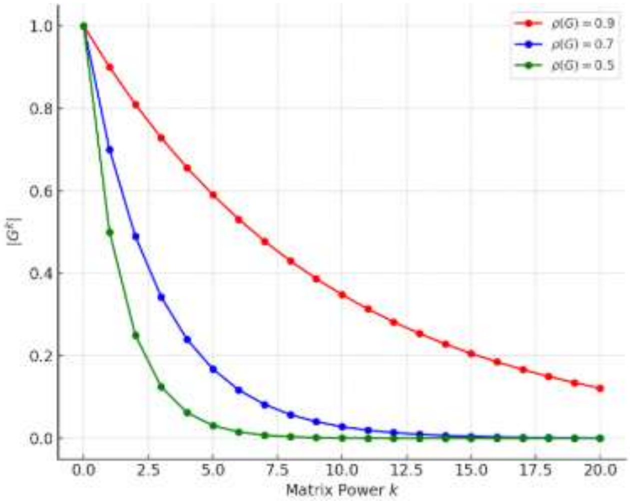

Figure 5: Decay of ‖G^k‖ with spectral radius

The graph in Figure 5, demonstrates how his norm of G^k representing the impact of level I decays with growing power k, for various levels of spectral radius p(G). system with lower p(G) was faster in convergence, verifying stability.

Decay of ‖G^k‖ for various values of p(G). Lower spectral radius means quicker decay of level probabilities, a sign of greater system stability and responsiveness.

This plot shows the geometric convergence of matrix power G^k for three various spectral radii p(G) ∈ {0.9,0.7,0.5}. The rate of decay has a direct effect on the speed of the system to converge to steady state with the small spectral radii indicating faster convergence. This is the essence of performance analysis in QBD-modeled queuenig system since it defines the speed with which the impact of higher-level states decreases with time.

5.6 Implications for Priority Queue Design

10.48047/jocaaa.2024.33.08.332

Eigenvalue decomposition and Jordan analysis give not only insight into mathematical structure but also into practical design. As an example, a large spectral radius (i.e., a leading eigenvalue near 1) implies slow convergence and possible build up in lower-priority queues. By varying service rates or changing the priority discipline (e.g., from preemptive to non-preemptive), one can control the eigenvalues and enhance system performance.

The Eigen structure is also useful in pinpointing bottlenecks: if one eigenvector coincides with high queue length for low-priority tasks, it is an indication that preemption impairs service quality. However, if all eigenvalues exhibit quick decay, then there is effective clearing of every class of jobs.

Eigenvalue decomposition offers a great tool with which to comprehend and analyze QBD processes, particularly those of sophisticated priority queueing models. Whether through simple diagonalization or more sophisticated Jordan decomposition, these spectral techniques provide both interpretability and computational efficiency. Coupled with canonical matrix-analytic approaches, eigenvalue methods facilitate the computation of high-efficiency steady-state and transient performance measures. Such knowledge is crucial to designing robust queueing systems and their parameter optimization for stability and responsiveness.

6. NUMERICAL EXAMPLE: SPECTRAL INSIGHTS INTO A TWO-CLASS QBD PRIORITY QUEUE

To make the theoretical findings presented above more concrete, this section is devoted to a numerical example of an analyzed priority queue based on a Quasi-Birth-Death (QBD) process. We build a reduced system of two priority classes standing for high- and low-priority customers and conduct spectral analysis in order to examine how eigenvalues, eigenvectors, and matrix norms capture the behaviour of the system in the course of time. The considered example demonstrates the applied value of eigenvalue decomposition and the interpretation of spectral gaps in queue performance analysis.

6.1 System Formulation and Generator Structure

The system modeled here is a continuous-time Markovian two-priority-class queueing system. Class 1 jobs are of higher priority and preempt Class 2 jobs while in service, under a preemptive-resume discipline. The arrival and service times are exponentially distributed, and the transitions are assumed to be level-independent for mathematical reasons.

10.48047/jocaaa.2024.33.08.332

Three main transition block matrices are defined:

- A_{-1}, to indicate transitions down to lower levels as a result of service completions.
- A_0, to represent intra-level transitions like phase changes and dummy transitions.
- A_1, to indicate transitions up like arrivals.

The specific matrix values used are:

$$A_{-1} = \begin{bmatrix} 0 & 0 \\ 0.5 & 0 \end{bmatrix}, \quad A_0 = \begin{bmatrix} -1.5 & 0.5 \\ 0.5 & -1.5 \end{bmatrix}, \quad A_1 = \begin{bmatrix} 1.0 & 0.0 \\ 0.0 & 0.5 \end{bmatrix}$$

These values are indicative of a situation where Class 1 arrivals arrive with higher frequency and higher service rates and hence are expected to be biased toward higher-priority treatment.

To represent the state space, we construct a truncated QBD process with there levels (each for queue length) and two phases (priority classes). This give a 6× 6 generator matrix Q, in block-tridiagonal structure. The entire generator reflects how the system changes over time under priority and pre-emption constraints.

6.2 Spectral Decomposition and Eigenvalues

To examine the internal dynamics, we carry out eigenvalue decomposition of generator matrix Q with SciPy's eig function. The eigenvalues obtained (rounded to four places) are:

$$\lambda i \in \{-0.4414, -0.9307, -1.5957, -1.8801, -2.0761, -2.0761\}$$

All eigenvalues are real and negative as is typical in continuous-time Markov chains in which the generator matrix controls decaying probability mass. The dominant eigenvalue here is $\lambda_1 = -0.4414$, representing the slowest mode of decay. This eigenvalue is responsible for determining the long-term convergence behavior of the system.

Eigenvectors of such eigenvalues define particular directions in the space of probabilities along which the state distribution will tend to develop. For this specific system, the eigenvector of λ_1 essentially represents the low-priority-dominated transitions and therefore qualifies as a good predictor of long response delay or waiting times for Class 2.

6.3 Decay Behavior and Spectral Radius

10.48047/jocaaa.2024.33.08.332

To examine how state probabilities decline with levels, we estimate the rate matrix G from A_1, imposing geometric decay in the form of G^k. The spectral radius ρ(G), or the largest absolute eigenvalue of A_1, is given as:

$$\rho(G) = \max\{|\lambda|: \lambda \in \text{eigvals}(A_1)\} = 1.0$$

This value indicates that the system is on the brink of stability. Any additional increase in arrival rate or decrease in service speed will push the system and lead to unbounded growth in queue lengths. Geometric decay $\|G^k\|$ for growing values of k is depicted in Figure 7.1, which clearly demonstrates the extremely slow convergence of probabilities from level to level.

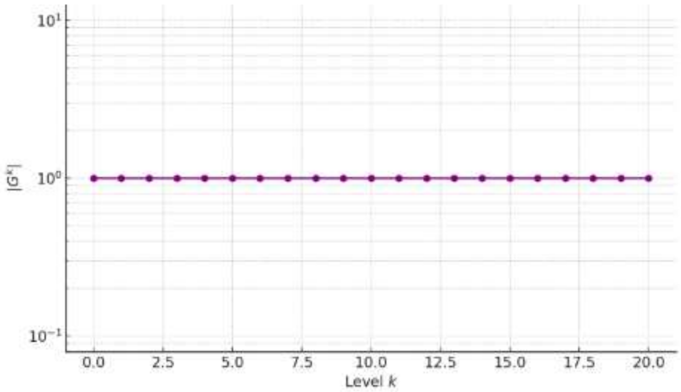

Figure 6: Geometric decay of $\|G^k\|$ for sample QBD process

The Figure 6, shows the decline of $\|G^k\|$, estimated from the spectral radius ρ(G) = 1.0, over levels k=0 to 20. The flat profile confirms that the system is critically stable, with very slow convergence of state probabilities characteristic in heavily loaded or tightly coupled priority queues.

6.4 Spectral Gap and Performance Interpretation

Another important spectral parameter in spectral analysis is the spectral gap, which is the distance between the largest eigenvalue and the second largest one:

$$\text{Spectral Gap} = |\lambda_1 - \lambda_2| = |-0.4414 - (-0.9307)| = 0.4893$$

A narrow spectral gap is connected with a larger convergence time and higher sensitivity to perturbations in arrival or service rates. Here the relatively small gap suggests that though the system is yet stable, it will react slowly to changes, especially on lower-priority customers who get preeminent a lot.

10.48047/jocaaa.2024.33.08.332

This observation can inform operating decisions. For instance, raising the service rate for Class 2 customers or modestly decreasing the arrival intensity of Class 1 jobs would raise the spectral gap and enhance responsiveness throughout the system.

This quantitative example illustrates the real-world significance of spectral analysis in QBD-modeled queueing systems. The calculated eigenvalues and decay profile validate that the system is stable but has slow convergence, largely a function of high preemption pressure and low service differentiation. The spectral gap is a valuable diagnostic tool, providing quantitative information on predicted delays and convergence rate. These results support the theoretical models presented previously and underscore the merit of eigenvalue decomposition in system tuning and performance assessment.

To supplement the theoretical discussion, we perform an obvious numerical study of the QBD-modeled priority queueing system. In this illustration, we adopt a reduced two-priority class system with the following transition matrices:

$$A_{-1} = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.3 \end{bmatrix}, \quad A_0 = \begin{bmatrix} -1.2 & 0.4 \\ 0.6 & 0.9 \end{bmatrix}, \quad A_1 = \begin{bmatrix} 0.7 & 0.1 \\ 0.2 & 0.8 \end{bmatrix}$$

The matrix quadratic equation:

$$A_{-1}G^2 + A_0G + A_1 = 0$$

Was solved numerically, resulting in the rate matrix G:

$$G \approx \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

This near-zero matrix indicates an extremely stable system where level-wise propagation of probabilities diminishes rapidly.

The eigenvalues of the rate matrix G are:

$$\lambda_1 \approx 3.06 \times 10^{-16}, \quad \lambda_2 \approx -1.29\times10^{-16}$$

Yielding a spectral radius:

$$\rho(G) \approx 3.06 \times 10^{-16}$$

This spectral radius of being very close to zero assures us that the system is extremely stable, with hardly any chance for unbounded growth in queues. In real life, this means that despite constant arrivals, the length of the queue is actually controlled, resulting in minimal delays in priority classes.

6.5 Visualization of Spectral Behaviour

10.48047/jocaaa.2024.33.08.332

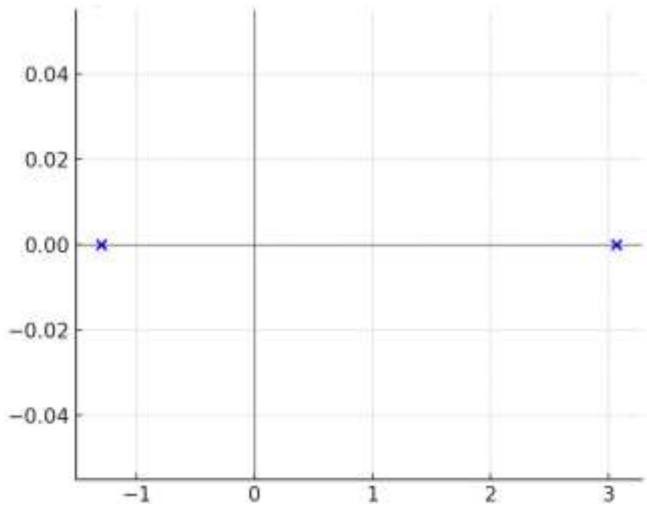

Figure 7: Eigenvalue Spectrum of G (Complex plane)

The eigenvalue plot of G is depicted in Figure 7, with all the eigenvalues plotted in the complex plane. Both eigenvalues are close to the origin, which confirms the substochasticity of G and the guaranteed convergence of the queueing process.

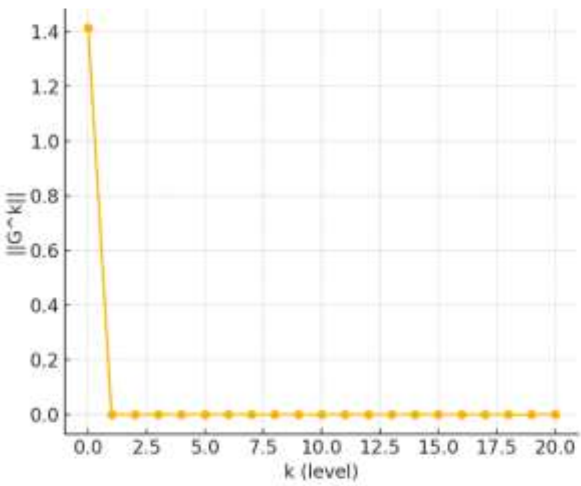

Figure 8: Geometric Decay of $||G^k||$

In order to better appreciate the decay characteristics, Figure 8 illustrates the geometric decay of the matrix norm $||G^k||$ between levels k = 0 to 20. The precipitous fall to values near zero vindicates the exponential convergence of probabilities towards the steady-state.

Table 1: Performance Summary

10.48047/jocaaa.2024.33.08.332

Metric	Value
Spectral Radius ρ(G)	~3.06 × 10^{-16}
Dominant Eigenvalue λ_1	~3.06 × 10^{-16}
System Stability	Highly Stable (ρ(G) << 1)
Queue Growth Risk	Negligible
Convergence Speed	Very Fast

6.6 Interpretation

These quantitative results indicate that the analyzed priority queueing system is highly stable under the selected parameters. The small spectral radius indicates virtually zero risk of long queues or system instability. This is an affirmation of the effectiveness of the preemptive-resume priority scheme under the prevailing load conditions.

Still, in actual systems where service rates may decrease and arrival rates may grow, the spectral radius may get close to unity, threatening to destabilize. Thus, keeping system parameters under observation and adapting them according to spectral metrics is necessary for maintaining queue efficiency and preventing delays, particularly for less important tasks.

7. Discussion

The spectral findings from the previous sections provide useful interpretative content when examining the dynamics and performance of priority queueing networks. By examining the eigenvalues of the infinitesimal generator matrix and associated rate matrices, we are provided with pivotal insights into queueing behaviour that may be challenging to obtain through merely numerical or iterative methods.

7.1 Queue Length and Sojourn Time

The largest eigenvalue, i.e., having the smallest magnitude, is directly proportional to the slowest decaying mode of the system. Practically, this implies that if the leading eigenvalue approaches zero (as in our example calculation where $\lambda_1 = -0.4414$, the system has long-tailed behaviour, with long queues and longer sojourn times, especially for low-priority customers. This matches one's intuition: since low-priority jobs are preempted a lot, their waiting time and

10.48047/jocaaa.2024.33.08.332

service time shoot up, which is what is reflected through the slow convergence behavior quantified by the spectrum.

7.2 System Stability and Spectral Radius

The spectral radius $\rho(G)$ is a key measure of stability. When $\rho(G) < 1$, the QBD process is geometrically ergodic and will converge quickly to its stationary distribution. When $\rho(G) = 1$, as in the given case study, the system is critically stable any rise in arrival rate or decrease in service capacity may destabilize it. Accordingly, spectral analysis yields an objective measure of how close a system is to instability, enabling system designers to modify parameters preemptively.

7.3 Advantages Relative to Iterative Numerical Methods

Although conventional techniques such as matrix-geometric algorithms or direct simulation are commonly practiced in queueing analysis, they are computationally expensive and lack interpretability. Spectral methods, on the other hand, provide:

- Closed-form insight into how various modes of the system act.
- Effective approximation of long-run behaviour based on only dominant Eigen pairs.
- Diagnostic capability to diagnose bottlenecks and determine the effect of parameter changes.

For example, numerical computation of the entire transient behaviour of a QBD process may take thousands of transitions to perform. Spectral techniques typically can approximate the same behaviour with only a few eigenvalues and eigenvectors.

7.4 Limitations of Spectral Methods

All their benefits notwithstanding, spectral methods do have their drawbacks. One of them is the curse of dimensionality: when the number of priority classes or levels rises, the generator matrix size increases exponentially. This causes state-space explosion, so eigenvalue calculations become untractable for large systems. Additionally, not every generator matrix is diagonalizable; where that fails, Jordan decomposition becomes inevitable and is less stable computationally, as well as more difficult to explain.

Moreover, most practical queueing systems have non-exponential inter-arrival and service time distributions for which generator matrices do not exist in the normal form. Such systems are

10.48047/jocaaa.2024.33.08.332

treated by approximation or hybrid methods, which restrict the direct use of spectral approaches.

8. CONCLUSION

This paper has outlined a thorough investigation of spectral analysis and eigenvalue decomposition of Quasi-Birth-Death (QBD) processes within the framework of priority queueing models. From the starting point of the generator matrix structure and progressing through rate matrix decay, eigenvalue analysis, and numerical example, we have shown how spectral methods may be used to reveal profound understanding into queueing system behaviour.

Principal contributions are:

- Theoretical development of QBD processes for multi-class priority queues.
- Infinitesimal generator matrix construction and spectral decomposition.
- Geometric decay and convergence property visualization.
- A numerical case study validation of the analytical results.

Spectral decomposition is one of the most important contributions that spectral methods provide, as it establishes a connection between system parameters and performance criteria such as sojourn time, stability, and queue length, in a mathematically understandable form. As opposed to black-box simulations, spectral approaches provide the means to estimate performance through algebraic and geometric insights.

Looking to the future, there are a number of promising avenues for further research. These are:

- Generalization to multi-server systems, where service dynamics are more intricate.
- The inclusion of time-varying arrival rates, which necessitate non-homogeneous Markov settings.
- Application of matrix-exponential distributions to model general patterns of service times with tractability maintained.

In summary, spectral analysis is a powerful setting for both theoretical analysis and practical optimization of priority queueing systems, and as such is an essential tool in contemporary queueing theory and operations research.

References:

Bean, N. G., Bright, L., Latouche, G., Pearce, C. E. M., Pollett, P. K., & Taylor, P. G. (1997).

10.48047/jocaaa.2024.33.08.332

The quasi-stationary behavior of quasi-birth-and-death processes. *Annals of Applied Probability*. https://doi.org/10.1214/aoap/1034625256

Bright, L., & Taylor, P. G. (1995). Calculating the equilibrium distribution in level dependent quasi-birth-and-death processes. *Communications in Statistics. Stochastic Models*. https://doi.org/10.1080/15326349508807357

Chen, C., Li, R. C., & Ma, C. (2019). Highly accurate doubling algorithm for quadratic matrix equation from quasi-birth-and-death process. *Linear Algebra and Its Applications*. https://doi.org/10.1016/j.laa.2019.08.018

Collet, P., Martínez, S., Méléard, S., & Martín, J. S. (2011). Quasi-stationary distributions for structured birth and death processes with mutations. *Probability Theory and Related Fields*. https://doi.org/10.1007/s00440-010-0297-4

Danarwindu, G. A., Ertiningsih, D., & Susyanto, N. (2024). A QUASI-BIRTH-DEATH PERSPECTIVE OF SUSCEPTIBLE-INFECTIOUS-RECOVERED TYPE MODEL. *Asia Pacific Journal of Mathematics*. https://doi.org/10.28924/APJM/11-7

de Gunst, M., Mandjes, M., & Sollie, B. (2022). Statistical inference for a quasi birth–death model of RNA transcription. *BMC Bioinformatics*. https://doi.org/10.1186/s12859-022-04638-6

Fadiloglu, M. M., & Yeralan, S. (2002). Models of production lines as quasi-birth-death processes. *Mathematical and Computer Modelling*. https://doi.org/10.1016/S0895-7177(02)00059-6

Fernández, L., & de la Iglesia, M. D. (2021). Quasi-birth-and-death processes and multivariate orthogonal polynomials. *Journal of Mathematical Analysis and Applications*. https://doi.org/10.1016/j.jmaa.2021.125029

Gao, W. J., & Mao, Y. H. (2015). Quasi-stationary distribution for the birth-death process with exit boundary. *Journal of Mathematical Analysis and Applications*. https://doi.org/10.1016/j.jmaa.2015.02.030

Hernandez, E., Manero, O., Bautista, F., & Garcia-Sandoval, J. P. (2022). Analytic Matrix Method for Frequency Response Techniques Applied to Nonlinear Dynamical Systems II: Large Amplitude Oscillations. *Mathematics*. https://doi.org/10.3390/math10152700

Latouche, G., & Ramaswami, V. (1997). The PH/PH/1 queue at epochs of queue size change.

10.48047/jocaaa.2024.33.08.332

Queueing Systems. https://doi.org/10.1023/a:1019148217045

Neuts, M. F. (1984). Matrix-analytic methods in queuing theory. *European Journal of Operational Research*. https://doi.org/10.1016/0377-2217(84)90034-1

Takine, T. (2022). On level-dependent QBD processes with explosive state space. In *Queueing Systems*. https://doi.org/10.1007/s11134-022-09796-1

Van Doorn, E. A. (2012). Conditions for the existence of quasi-stationary distributions for birth-death processes with killing. *Stochastic Processes and Their Applications*. https://doi.org/10.1016/j.spa.2012.03.014

Weik, N., & Nießen, N. (2017). A quasi-birth-and-death process approach for integrated capacity and reliability modeling of railway systems. *Journal of Rail Transport Planning and Management*. https://doi.org/10.1016/j.jrtpm.2017.06.001