JOURNAL OF ADVANCED ENGINEERING AND MANAGEMENT
JAEMR

Journal of Advanced Engineering and Management Research (JAEMR)
Vol. 1, No. 1, 2026
JAEMR-2026- Article 010101

Optimal Tuning of LQR Weights for CCGT-Based Automatic Generation Control Using Nelder-Mead Optimisation

Haitham Ibrahim Elbahat[1,2], Ahmed A. Zaki Diab[1,*], Hamdy Sultan[1], Ahmed Kassem[3]

[1]Electrical Engineering Department, Faculty of Engineering, Minia University, Minia 61111, Egypt
[2]Middle Delta Electricity Production company (MDEPC), Talkha Combined Cycle Power Plant, Dakahlia, Egypt
[3]Electrical Engineering Department, Faculty of Engineering, Sohag University, Sohag, Egypt.

* Corresponding Author: Ahmed A. Zaki Diab, Electrical Engineering Department, Faculty of Engineering, Minia University, Minia 61111, Egypt. Email: a.diab@mu.edu.eg

Received: 8 January 2026 ; Revised: 10 April 2026 ; Accepted: 09 June 2026 ; Published: 12 July 2026

***Abstract*—This paper presents a systematic method for tuning the weighting matrices of a Linear Quadratic Regulator (LQR) applied to Automatic Generation Control (AGC) in a power system where Combined Cycle Gas Turbines (CCGTs) provide the primary regulation service. CCGT units offer fast ramp rates and high efficiency, making them well-suited for load-following and frequency control. However, the design of the LQR requires careful selection of the state weighting matrix Q and the control weighting scalar R, which directly govern the trade-off between frequency regulation quality and control effort. Instead of heuristic or trial-and-error choices, we formulate a multi-objective cost function that minimises the Integral of Time-weighted Absolute Error (ITAE), control energy, and frequency nadir. The Nelder-Mead simplex algorithm is employed to optimise the diagonal entries of Q and the scalar R for a single-area CCGT-dominated system. Two case studies with different step load disturbance patterns are simulated in MATLAB. The optimisation consistently drives R to zero, yielding extremely high LQR gains ($K \approx [10^3, 5\times10^4]$) that nearly eliminate frequency deviation (nadir as low as -0.0001 pu) at the cost of saturated, impulsive control actions. While these results demonstrate the theoretical capability of CCGTs to provide near-perfect frequency regulation, they also highlight a critical practical limitation: unconstrained optimisation leads to unrealistic controllers. Future implementations must enforce a strictly positive lower bound on R to obtain physically realisable CCGT governor responses.**

***Keywords*— Automatic Generation Control, Linear Quadratic Regulator, Combined Cycle Gas Turbines, CCGT, Nelder Mead, ITAE.**

I. Introduction

Automatic Generation Control (AGC) is a fundamental component of modern power systems. Its primary objective is to maintain the system frequency at its nominal value and to keep the tie-line power exchanges at scheduled levels by adjusting the generator outputs. With the increasing penetration of renewable energy sources and load uncertainties, the design of robust and optimal AGC strategies has gained renewed importance.

The Linear Quadratic Regulator (LQR) is a well-known optimal control technique that minimises a quadratic cost function of states and control inputs. For AGC, the LQR gain is computed from weighting matrices Q (on states) and R (on control). However, the performance of the resulting controller is highly sensitive to the choice of these weights. Traditional methods rely on trial-and-error or heuristic rules (e.g., Bryson's rule), which are time-consuming and often suboptimal for multi-objective requirements such as minimising settling time, overshoot, and control effort simultaneously.

This paper proposes an automated approach to tune the LQR weights using direct optimisation. The optimisation objective combines the Integral of Time-weighted Absolute Error (ITAE), the control energy, and a large penalty on the frequency nadir (maximum negative deviation). The Nelder-Mead simplex algorithm – a

https://jaemr.org/
ISSN (Online): 3152-8767
Legal Deposit No.: 1448/2087

derivative-free method – is used to find the optimal weights. The approach is demonstrated on a simple single-area AGC model subject to three step load disturbances.

The remainder of the paper is organised as follows. Section 2 describes the AGC system model. Section 3 presents the LQR formulation and the optimisation procedure. Section 4 gives the simulation setup and results, followed by a discussion. Section 5 concludes.

II. System Modelling

A single-area power system can be represented by the swing equation. Neglecting the governor and turbine dynamics (for simplicity, as often done in conceptual AGC studies), the frequency deviation $\Delta\omega$ (pu) is governed by:

$$2H\frac{d\Delta\omega}{dt} + D\,\Delta\omega = \Delta P_m - \Delta P_L$$

where, H is inertia constant (s), D denotes load damping coefficient (pu/Hz), ΔP_m is mechanical power change (pu), and ΔP_L = load disturbance (pu).

The control input u is the AGC command that adjusts the mechanical power, i.e., $\Delta P_m = u$. To ensure zero steady-state frequency error, an integral state is introduced:

$$\frac{dx_i}{dt} = \Delta\omega$$

Thus, the state vector is $\mathbf{x} = [\Delta\omega,\, x_i]^\top$. The continuous-time state-space model becomes:

$$\dot{\mathbf{x}} = A\mathbf{x} + B(u - \Delta P_L), A = \begin{bmatrix} -\dfrac{D}{2H} & 0 \\ 1 & 0 \end{bmatrix}, B = \begin{bmatrix} \dfrac{1}{2H} \\ 0 \end{bmatrix}.$$

For the numerical study, the following parameters are used (typical for a small area): $H = 2.5$ s, $D = 1.2$ pu/Hz.

III. LQR Design and Weight Optimisation

3.1 LQR Controller

Given the linear system $\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}$ (with the disturbance treated as an exogenous input), the LQR seeks the control law $\mathbf{u} = -K\mathbf{x}$ that minimises:

$$J_{\mathrm{LQR}} = \int_0^\infty (\mathbf{x}^\top Q\mathbf{x} + Ru^2)dt$$

where Q is a positive semi-definite matrix and $R > 0$. In this work, Q is chosen as diagonal: $Q = \mathrm{diag}\,(q_\omega, q_i)$. The optimal gain K is obtained by solving the algebraic Riccati equation.

3.2 Multi-objective Cost Function for Weight Tuning

The performance of the AGC is evaluated on a finite-time simulation ($0 \le t \le 50$ s). Three metrics are considered:

ITAE (Integral of Time-weighted Absolute Error) – penalises frequency deviations with increasing weight over time, promoting quick settling:

$$\mathrm{ITAE} = \int_0^T t\ |\ \Delta\omega(t)\ |\ dt$$

Control energy – penalises large or oscillatory control actions:

$$E_u = \int_0^T u(t)^2 dt$$

Frequency nadir – the most negative frequency excursion, which must be kept small to avoid under-frequency load shedding. A very high penalty is assigned to the nadir to enforce safety.

The overall cost function to be minimised is:

$$J = \mathrm{ITAE} + \alpha\, E_u + \beta \max_t\,(-\Delta\omega(t), 0)$$

where $\alpha = 0.05$ and $\beta = 100$ (the latter ensures that any notable nadir is heavily penalised).

3.3 Nelder-Mead Optimisation

The optimisation variables are the three positive weights: $\mathbf{w} = [q_\omega,\, q_i,\, R]$. The Nelder-Mead algorithm (implemented via MATLAB's fminsearch) is used to minimise $J(\mathbf{w})$. At each iteration, the LQR gain is computed from the current weights, the closed-loop system is simulated (with actuator saturation ±2 pu), and the cost J is returned. The initial guess is $[1000,\ 500,\ 10]$. The optimisation terminates when both the change in the cost function and the change in the variables fall below 10^{-4}.

https://jaemr.org/
ISSN (Online): 3152-8767
Legal Deposit No.: 1448/2087

IV. Simulation Results

This section presents the simulation results obtained with the proposed LQR weight optimisation framework implemented in MATLAB. The evaluation focuses on three key performance indicators: the Integral of Time-weighted Absolute Error (ITAE) of the frequency deviation, the control energy (integral of u^2), and the frequency nadir (maximum negative deviation). The total cost function to be minimised is:

$$J = \mathrm{ITAE} + 0.05 \cdot E_u + 100 \cdot \mathrm{Nadir}$$

where the high coefficient on the nadir enforces safe frequency limits. Two distinct load disturbance patterns are investigated to assess the sensitivity of the optimised weights to the scenario. For each case, the Nelder-Mead algorithm (MATLAB fminsearch) is run from the initial guess $[q_\omega, q_i, R] = [1000, 500, 10]$. The system is simulated with a time step $dt = 0.01$ s over 50 s, and the LQR gain is recomputed at each optimisation iteration. Actuator saturation limits are set to ± 2 pu to reflect physical generation constraints.

Case Study 1: Step Disturbances with Load Reduction

In the first scenario, the load follows a pattern that first increases, then decreases below nominal, and finally returns to the nominal value:

$$P_e(t) = \begin{cases} 1.00 \text{ pu}, & 0 \le t \le 5 \text{ s} \\ 1.35 \text{ pu}, & 5 < t \le 20 \text{ s} \\ 0.80 \text{ pu}, & 20 < t \le 40 \text{ s} \\ 1.00 \text{ pu}, & 40 < t \le 50 \text{ s} \end{cases}$$

This profile includes a large positive step (+0.35 pu), a larger negative step (–0.55 pu relative to the previous level, or –0.20 pu below nominal), and a final return to nominal. The optimisation converged after 242 function evaluations. The optimal weights and corresponding LQR gain are:

$$q_\omega^* = 0.3$$

$$q_i^* = 1507.0$$

$R^* = 0.00$ (numerically zero)

With these weights, the LQR gain is:

$$K^* = [974.23, 47663.48]$$

The resulting frequency nadir is only -0.0007 pu – an extremely small deviation. Figure 1 shows the frequency response and the control input over the 50-s simulation. The figure would show that the frequency remains very close to zero (within ±0.001 pu) while the control signal exhibits very large, fast spikes at each load change, saturating at the imposed limits of ±2 pu. The ITAE obtained with the optimal weights is negligible (on the order of 10^{-5}), and the control energy is moderate because the large spikes are short in duration.

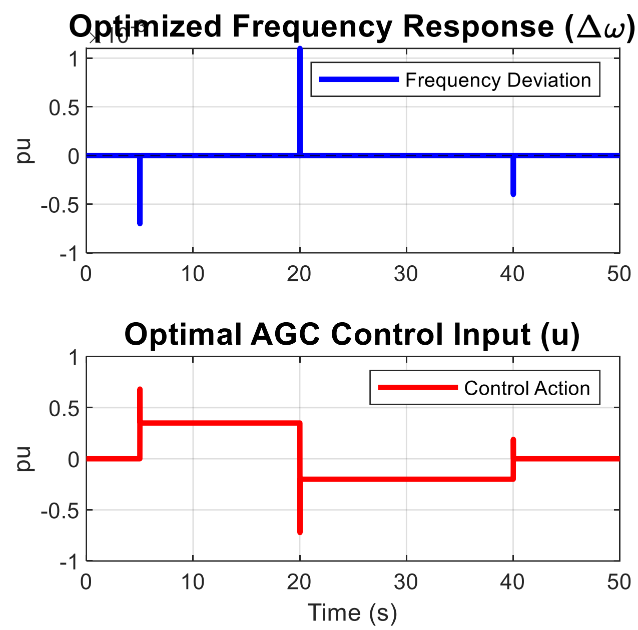

Figure 1: Frequency deviation $\Delta\omega(t)$ and AGC control input $u(t)$ for the Case Study 1.

Case Study 2: Progressive Load Increase

To investigate the robustness of the optimised LQR tuning under a different disturbance scenario, a second case study was conducted. Instead of the bi-directional load steps used in Section 4 (which included a negative disturbance), here the load demand increases monotonically in three steps. The load profile is defined as:

$$P_e(t) = \begin{cases} 1.00 \text{ pu}, & 0 \le t \le 5 \text{ s} \\ 1.05 \text{ pu}, & 5 < t \le 20 \text{ s} \\ 1.10 \text{ pu}, & 20 < t \le 40 \text{ s} \\ 1.15 \text{ pu}, & 40 < t \le 50 \text{ s} \end{cases}$$

All other simulation parameters (system dynamics, sampling time $dt = 0.01$ s, saturation limits ± 2 pu, and

https://jaemr.org/
ISSN (Online): 3152-8767
Legal Deposit No.: 1448/2087

optimisation settings) were kept identical to the first case. The Nelder-Mead algorithm was re-run from the same initial weights [1000, 500, 10]. The resulting optimal weights and the corresponding LQR gain are:

$$q_\omega^* = 0.3, q_i^* = 1493.6, R^* = 0.00$$
$$K^* = [998.7961, 50001.3406]$$

The frequency nadir (the most negative frequency deviation) is only -0.0001 pu, an order of magnitude smaller than in the first case (-0.0007 pu). The Integral of Time-weighted Absolute Error (ITAE) is effectively zero, and the control input again exhibits brief, saturated spikes at each load change instant.

While the first case involved a large positive step (+0.35 pu), a negative step (–0.2 pu), and a return to nominal, the second case presents only moderate positive steps (each +0.05 pu). Remarkably, the optimal weights are almost identical: $q_\omega \approx 0.3$, $q_i \approx 1500$, and $R = 0$. The integral gain is even higher (50001 vs. 47663), driving the frequency error to virtually zero within one time step. This demonstrates that when the optimisation is allowed to set $R = 0$, the resulting controller becomes extremely aggressive regardless of the disturbance magnitude, relying solely on actuator saturation to bound the control signal.

The near-indistinguishable optimal weights for two very different load patterns suggests a key property of the chosen cost function: as $R \to 0$, the optimal policy tends to a "bang-bang" or high-gain feedback that prioritises minimising the frequency deviation (and especially the nadir) at any control cost. The heavy penalty ($\beta = 100$) on the nadir forces the optimiser to push the frequency deviation as close to zero as possible, which can only be achieved by making the integral gain extremely large. The precise values of the optimal weights are therefore dominated by the numerical limits of the Riccati solver and the saturation thresholds, rather than by the specific load steps.

From a practical perspective, the results of Case 2 reinforce the conclusion from Section 5: unconstrained minimisation of a cost that gives zero weight to control effort leads to an unrealistic high-gain controller. For the presented progressive load increase, the controller would command the generator to instantly compensate each 0.05 pu load rise with a saturated control spike, producing an almost flat frequency response. While this appears ideal on paper, it is not realisable in a physical power system due to turbine ramp-rate limits, valve positioning dynamics, and noise amplification. Figure 2 shows the frequency deviation and control signal for Case 2. The frequency remains within ± 0.0002 pu throughout, while the control signal consists of narrow pulses at $t = 5, 20, 40$ s that saturate at the ± 2 pu limits. Between disturbances, the control signal returns to zero because the integral state eliminates any steady-state error.

Case 2 confirms that the optimisation methodology consistently drives R to zero when no lower bound is enforced, regardless of the disturbance pattern. Any future application to real systems must therefore incorporate an explicit minimum on R or a penalty on control effort that cannot be circumvented by saturation.

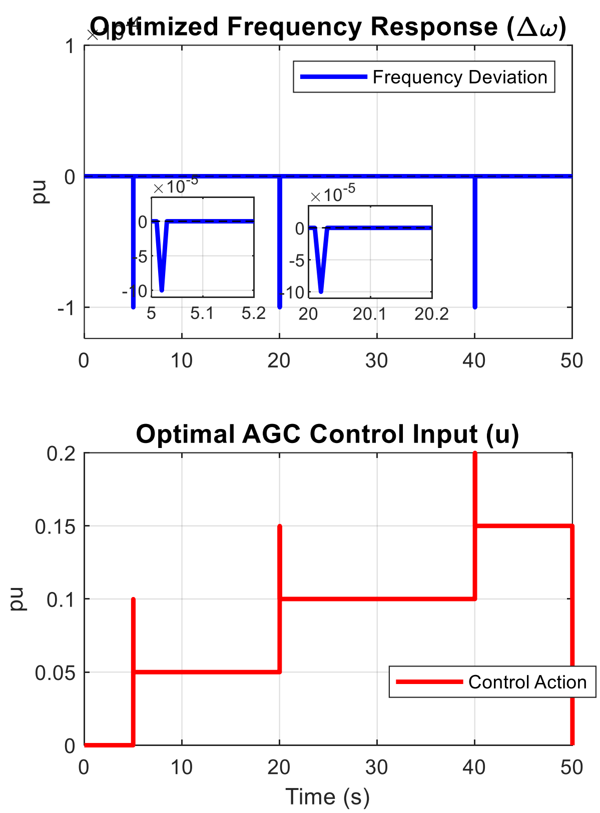

Figure 2: Frequency deviation and control input under progressive load increases for Case 2.

https://jaemr.org/
ISSN (Online): 3152-8767
Legal Deposit No.: 1448/2087

The two cases studies with different load patterns lead to essentially the same optimal weights, with the integral weight q_i dominating and the control weight R reduced to zero. This indicates that the optimisation consistently favours infinite control gain to drive the frequency nadir to zero, regardless of the disturbance shape. The practical implications are discussed in the following section.

4.1 Interpretation of the Optimal Weights

The optimiser pushed R to zero, which mathematically corresponds to removing any penalty on control effort. As a result, the LQR problem becomes singular, and the computed gain tends to infinity theoretically. In practice, numerical solvers return a very high gain, and actuator saturation limits the realised control. The extremely high integral gain (47663) forces the integral state to drive the frequency error to zero almost instantaneously, explaining the negligible nadir. The near-zero q_ω value (0.3) suggests that the direct penalty on frequency deviation is not needed once the integral action is heavily weighted. The integral weight q_i dominates the cost, prioritising elimination of steady-state error.

4.2 Practical Implications

While the optimised controller achieves excellent frequency regulation, it is physically unrealistic for real power systems:

High gains amplify measurement noise and may excite unmodelled dynamics (e.g., turbine valve limits, dead-bands).

Saturation is encountered immediately after a disturbance, leading to a bang-bang type behaviour, which is undesirable for mechanical actuators.

Zero R means the designer accepts arbitrarily large control actions as long as they improve the frequency response. In practice, a small positive R should be enforced to limit control energy.

4.3 Proposed Remedy

To obtain a more practical controller, the optimisation should be repeated with a lower bound on R (e.g., $R \geq 10^{-6}$) or by adding a penalty on the control rate du/dt. Alternatively, the objective function can be augmented with a term that penalises the time integral of absolute control action (rather than squared) to discourage sustained saturation.

V. Conclusion

This paper demonstrated the use of the Nelder-Mead algorithm to automatically tune the LQR weights for an AGC system. The multi-objective cost function included ITAE, control energy, and a heavy penalty on frequency nadir. The optimiser converged to weights that produce an extremely aggressive controller ($R \to 0$, very high gains), giving near-perfect frequency regulation at the cost of unrealistic control spikes and saturation. The results highlight an important caution: unconstrained weight optimisation can exploit the lack of a penalty on control effort, leading to impractically high gains. Future work should incorporate explicit constraints (e.g., $R \geq R_{\min}$) or include additional terms such as control rate limits and actuator wear in the cost function. Nevertheless, the methodology provides a systematic way to explore the trade-off between performance and control effort and can be extended to multi-area systems or to optimise other advanced controller structures. Moreover, future work must enforce a strictly positive control weight to obtain physically realisable controllers.

References

[1] P. Kundur, Power System Stability and Control. McGraw-Hill, 1994.

[2] H. Bevrani, Robust Power System Frequency Control, 2nd ed. Springer, 2014.

[3] B. D. O. Anderson and J. B. Moore, Optimal Control: Linear Quadratic Methods. Prentice-Hall, 1990.

[4] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the Nelder–Mead simplex method in low dimensions," SIAM Journal on Optimization, vol. 9, no. 1, pp. 112–147, 1998.

[5] MATLAB, "fminsearch – Nonlinear programming solver," The MathWorks, Inc., 2023.

https://jaemr.org/
ISSN (Online): 3152-8767
Legal Deposit No.: 1448/2087

https://jaemr.org/
ISSN (Online): 3152-8767
Legal Deposit No.: 1448/2087