

Optimizing Electricity Distribution in Power Grid: A Graph Theory and Reinforcement Learning Framework

Ziqi Zheng

lisa.zheng000@gmail.com

Maria Carrillo High School, Nontrivial Fellowship Program

Keywords: energy optimization, graph theory, reinforcement learning, Markov Decision Processes, Ford-Fulkerson, Dijkstra's, Q-Learning

Abstract

In this paper two kinds of algorithms are proposed to improve the power grid distribution in which one uses static methods (Dijkstra's, Ford-Fulkerson) to consider the capacity/loss from plant and transmission and another is probability/reinforcement learning based method (Markov Decision Processes, and Q-Learning) to take into consideration the uncertainties, such as fuel shortages and wind variability for the goal of optimizing its energy flow. We took the dataset for Cuba's power plants as a case study to test the effectiveness of these algorithms for power distribution. Our results show a major potential improvement of 22-68% in energy generation from using these two kinds of algorithms (static/probability) compared to the current country's available operating power.

Background

We chose Cuba to specifically study due to availability of its data on energy capacities and losses, but these algorithms can be applied to other countries. Cuba's electricity system is obsolete and in urgent need of repair. On an average day, the Cuban national government can only meet 50-70% of the country's needs. The National Electric System, built after 1959, has not received consistent maintenance for 35 years. Oil-based power plants have suffered an energy crisis due to dwindling imports from Mexico, Russia, and Venezuela. 10%, or 1 million people, left Cuba between 2022 and 2023, the largest migration in the country's history. The Cuban government warned that it may have to further increase gas and electricity prices to encourage its citizens to cut back on energy usage. Cuba's power plants also suffer because 85% of the plants heavily rely on poor-quality crude oil that is corrosive due to its high amount of sulfur, exacerbating the deterioration of boilers, turbines, and pipes.

The capital of Cuba, Havana, requires a minimum of 3000 MW of energy for its citizens, but the island's energy grid can only generate 1700 MW-1900 MW on average. In this paper, we will use optimizing methods, both static and dynamic to test the improvement efficiency. We got the simulated results of 2700-2900 MW of energy (20%-68% increase) on average based on the algorithms. Specifically, Dijkstra's results in 2325.25 MW, Ford-Fulkerson in 2800.00 MW, Markov Decision Process in 3198.38 MW, and Q-Learning in 2739.31 MW.

Method

We use NetworkX/DiGraph (nodes/edges) to simulate the electric grid with nodes representing power plants and edges with attributes of the capacity and loss of the transmission line. As shown in Figure 1, 17 major power plants generate power to send to one of two substations, which then direct to the final destination, the main Havana load. The real-time capacities of the power plants are listed in Table 1. The total delivered power to Havana is calculated as follows:

$$\text{Power delivered for Havana Load} = \sum C_i * (1 - L_i)$$

where

$C_i = \min(\text{power capacity of power plant } i, \text{ capacity of transmission line from plant } i \text{ to its substation})$

$L_i = \text{power loss of transmission line from plant } i \text{ to its substation}$

Figure 1: Equation for total power delivered

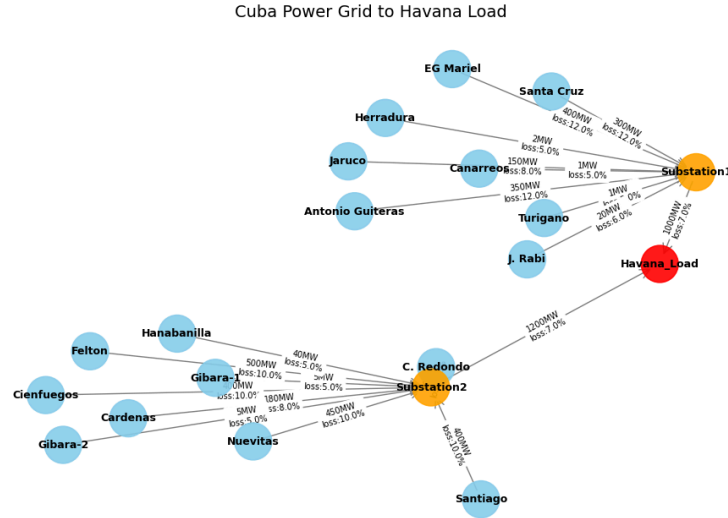


Figure 2: connected graph of power plants with two substations

Fuel power plant	Power Plant capacity	Fuel power plant	Power Plant capacity	Fuel power plant	Power Plant capacity
EG Mariel	450	Santiago	450	Gibara-1	5.1
Termoeléctrica de Santa Cruz	300	Jaruco	163	Gibara-2	4.5
Antonio Guiteras	330	Cardenas	193	J. Rabi	20
Cienfuegos	382	Canarreos	0.27	C. Redondo	55
Nuevitas	439	Turiguano	0.45	Hanabanilla dam	45
Felton	500	Herradura	1.5		

Table 1: energy capacities for power plants in Cuba

Algorithms

1. Dijkstra's Algorithm

From the source node to the sink (Havana Load), Dijkstra's (the shortest-path algorithm in graph theory) iterates through paths with the goal of minimizing the total sum of the losses along transmission lines. The weight on the edge in this case is the attribute "loss". When the shortest path is determined, the total delivered power is calculated

using Figure 1. However, this may not be the totally optimal result because Dijkstra's does not consider the capacities and losses together, only on minimizing the loss.

```
for plant in plants.keys():
    shortest_path = nx.dijkstra_path(G, plant, 'Havana_Load', weight='loss')
    path_loss = sum(G[u][v]['loss'] for u, v in zip(shortest_path[:-1], shortest_path[1:]))
    delivered = min(G.nodes[plant]['capacity'], G[shortest_path[0]][shortest_path[1]]['capacity'])
    delivered_power = delivered * (1-path_loss)
    results.append({
        'Plant': plant,
        'Path': shortest_path,
        'Total Loss (%)': path_loss * 100,
        'Delivered Power (MW)': delivered_power
    })
```

2. Ford-Fulkerson

Ford-Fulkerson requires an external source node as the start; it then traverses through the graph with the goal of maximizing the total weight on the edges, or the attribute “capacity” in this case. Each edge can be considered only at most once. However, this may not be the totally optimal result because Ford-Fulkerson does not consider the capacities and losses together, only on maximizing the capacity. Here, we set an average transmission loss of 15%, approximating Cuba's 15.29% from World Bank data. Shown in Figure 2, the super source node as a virtual node is connected to all plant nodes.

```
for plant in plants.keys():
    flow_value, flow_dict = nx.maximum_flow(G, plant, 'Havana_Load', capacity = 'capacity')
    delivered = min(flow_value, G.nodes[plant]['capacity'])
    delivered_power = delivered * (1-average_loss)
    results.append({
        'Plant': plant,
        'Flow Value (MW)': flow_value,
        'Delivered Power (MW)': delivered_power
    })
```

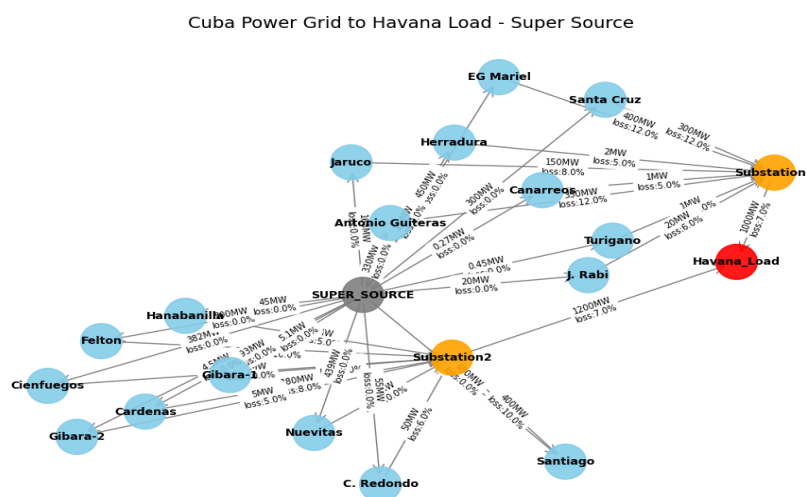


Figure 3: connected graph with super source

3. Markov Decision Processes (MDPs)

MDPs are used to handle uncertainty, such as grid failure, fuel shortage, or wind variability for wind farms. The virtual agent travels through a Markov Chain where each action, such as “reroute”, “do nothing”, or “take a load off” has a specific transition probability and receives rewards (positive for a desired action and negative for an undesirable action). This algorithm is used for the goal of learning the optimized actions to maximize the cumulative reward over iterations. Since using 17 plants requires too much computational power, 5 representative plants are used. As shown in Figure 4, the Bellman equation is used to recursively calculate the expected cumulated reward, using the gamma constant as a discount factor to prioritize short-term actions over long-term ones.

$$V(s) = \max_a (R(s, a) + \gamma V(s'))$$

where

$V(s)$ = the expected return value at current state s

\max_a = the maximum value of the action a

$R(s, a)$ = the reward for taking action a at state s

$\gamma V(s')$ = the value of the next state multiplied by the discount factor

Figure 4: The Bellman equation

Below is a snippet of the possible next states if the action “reroute” is taken.

```
if action.startswith('reroute'):
    #Prioritize the chosen plant
    priority_plant = 'Cienfuegos' if action == 'reroute_cienfuegos' else 'Santiago'
    path = nx.dijkstra_path(G, priority_plant, 'Havana_Load', weight = 'loss')
    path_loss = sum(G[u][v]['loss'] for u, v in zip(path[:-1], path[1:]))
    if priority_plant == 'Felton':
        priority_cap = 350 if state[0] == 'shortage' else 500
    elif priority_plant == 'Cienfuegos':
        priority_cap = 382
    elif priority_plant == 'Santiago':
        priority_cap = 280 if state[2] == 'shortage' else 405
    elif priority_plant == 'Gibara-1':
        priority_cap = 4.1 if state[3] == 'low' else 5.1
    elif priority_plant == 'Gibara-2':
        priority_cap = 3.6 if state[4] == 'low' else 4.5
    flow = min(priority_cap, G[path[0]][path[1]]['capacity'])
    total_flow += flow * (1 - path_loss)
    remaining_cap = min(1200 - flow, 1200)
```

4. Q-Learning

Q-learning is a method of reinforcement learning that learns the optimal policy within the environment of power plants with actions without using the models in MDPs but through the ϵ -greedy method of exploitation v. exploration. The Bellman equation adapted for Q-Learning shown in Figure 5 is used to calculate the cumulative Q-value. The Q-table informs the agent of which decision at each power plant leads to overall maximized reward and its values are updated as the computer is trained.

$$Q(S, A) \leftarrow Q(S, A) + \alpha(R + \gamma Q(S', A') - Q(S, A))$$

where

$Q(S, A)$ = current Q value for taking action A at state S
 α = the learning rate as a real number between 0 and 1
 R = the reward
 $\gamma Q(S', A')$ = the discount factor multiplied by the Q-value of taking next action A' at the next state S'

Figure 5: the Bellman equation for Q-Learning

```

for episode in range(episodes):
    for state in states:
        state_visits[state] += 1
        for _ in range(10):
            if random.random() < epsilon:
                action = random.choice(actions)
            else:
                action = max(actions, key=lambda a: Q[(state, a)])
            total_flow, reward = calculate_flow(state, action)

            next_state = simulate_grid_state()
            best_next_action = max(actions, key=lambda a: Q[(next_state, a)])
            Q[(state, action)] += alpha * (reward + gamma * Q[(next_state, best_next_action)] - Q[(state,
action)])

        epsilon = max(0.01, epsilon * 0.995)

```

Experimental Results

Delivered Power in Havana Load			
Dijkstra's	Ford-Fulkerson	MDPs	Q-Learning
2325.25 MW	2800.00 MW	3198.38 MW	2739.31 MW
2325.25/1900 = +22.38%	2800/1900 = +47.37%	3198.38/1900 = +68.34	2739.31/1900 = +44.17%

Table 2: results of delivered power for each algorithm

As shown in Table 2, these static and dynamic algorithms result in a 22-68% improvement in delivered energy compared to Cuba's current output.




Figure 6: delivered power per plant for Dijkstra's




Figure 7: delivered power per plant for Ford-Fulkerson

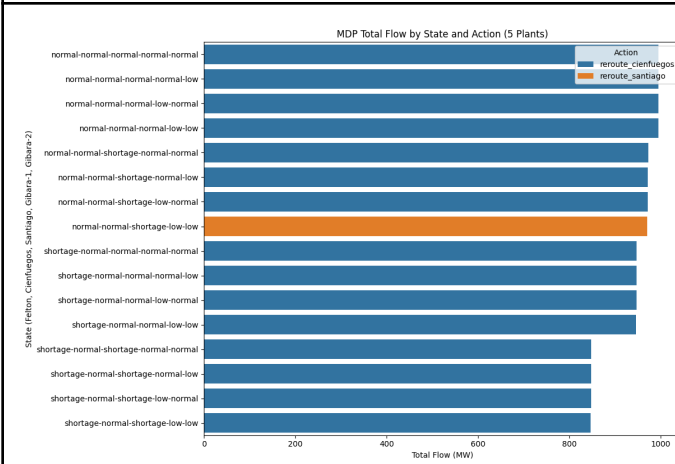


Figure 8: results for Markov Decision Processes

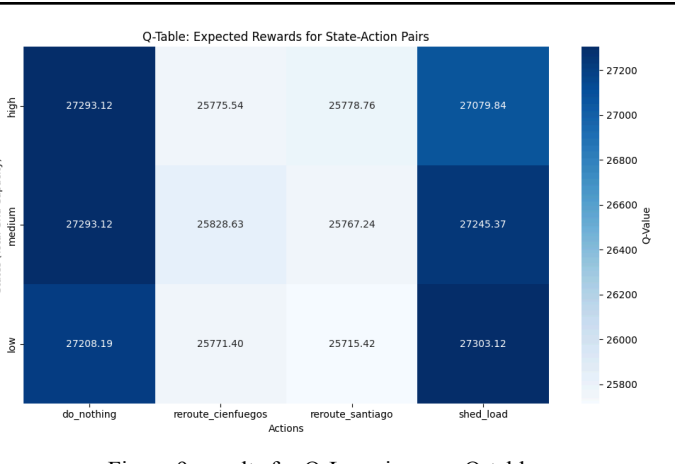


Figure 9: results for Q-Learning as a Q-table

Figure 6 and 7 show the delivered power per plant in the optimized route when Dijkstra's and Ford-Fulkerson are applied, respectively. Figure 8 shows the total power flow by state and action, and Figure 9 shows the Q-values for each action and state pair after many iterations, from which the optimal combinations can be deduced.

Conclusions

Dijkstra's Algorithm: This static algorithm that aims to minimize losses without considering variability delivers around 2325.25 MW, a 390 MW improvement compared to Cuba's baseline of 1935 MW. However, it falls short of Havana's 3000 MW demand, resulting in a high blackout risk (30-40%) due to its inability to adapt to fuel shortages and wind variability.

Ford-Fulkerson: Aiming to maximize flow through the grid with two substations (Substation 1: 1000 MW, Substation 2: 1200 MW), this algorithm delivers 2800 MW, a 474.74 increase from Dijkstra's. It reduces blackout frequency to 15-20% but also lacks dynamic adaptation to real-time uncertainties, compromising its reliability during fuel shortages.

MDPs: An extrapolated energy delivery from a subset of 5 plants delivers 3198.38 MW, exceeding Havana's demand. This algorithm assumes unpredictability and adapts accordingly to uncertainties via probabilistic

transitions. The blackout rate is reduced to <5%. However, this extrapolation assumes linear scaling, which could be overestimating the energy grid's true capabilities.

Q-Learning: The Q-Learning algorithm, which learns optimal actions based on rewards (do nothing, reroute to Cienfuegos, reroute to Santiago, and shed load) across high (>2500 MW), medium (2000-2500 MW), and low (<2000 MW) states, delivers 2739.31 MW. By shedding load (2400 MW) in low states 5-10% of the time, this algorithm reduces blackout frequency by ~20%, meeting the initial goal. The ability of this algorithm to make decisions based on optimal outcome and respond to dynamic conditions renders it highly effective.

Conclusion: Q-Learning and MDPs have the greatest potential for blackout reduction, reducing the rate to 10-15% and <5%, respectively. Q-Learning is more practical due to its direct incorporation of all 17 power plants. The 261-1000 MW gap from 3000 MW may be due to unrecoverable physical constraints, such as 15% losses and substation bottlenecks, but the overall 20% blackout reduction would significantly improve energy reliability for 2.1 million of Cuba's citizens, improve quality of life, and support services.

Google Colab link:

https://colab.research.google.com/drive/1LsQT_PrXtqRb_azo9Q6Clr87tTWS7gYE#scrollTo=2bQ_sM_aSt3n

References

1. Angulo, Radio. "Complex Solutions for Electricity Supply in Cuba - Radio Angulo." *Radio Angulo*, 22 May 2025, www.radioangulo.cu/en/2025/05/21/complex-solutions-for-electricity-supply-in-cuba/. "Cuba in Crisis as Electrical Grid Collapses Causing Island-Wide Blackout." *World Socialist Web Site*, 22 Oct. 2024, www.wsws.org/en/articles/2024/10/22/ckjs-o22.html.
2. Frank, Marc. "Cuba Warns of Worsening Blackouts as Fuel Crisis Bites." *Reuters*, 28 Sept. 2023, www.reuters.com/world/americas/cuba-warns-worsening-blackouts-fuel-crisis-bites-2023-09-28/.
3. Keenan, Lauren, and Jake Larsen. "Millions in Cuba without Power after Electrical Grid Collapse." *Straight Arrow News*, 4 Dec. 2024, san.com/cc/millions-in-cuba-without-power-after-electrical-grid-collapse/.
4. Nora Gámez Torres. "Cuba Admits to Massive Emigration Wave: A Million People Left in Two Years amid Crisis." *Miami Herald*, 19 July 2024, www.miamiherald.com/news/nation-world/world/americas/cuba/article290249799.html.
5. Salinas, Zoe. "Cuban Government's Excuse: 'No Single Solar Park Can Solve the Energy Deficit.'"" *CubaHeadlines*, 6 Apr. 2025, www.cubaheadlines.com/articles/300347.
6. Sanchez, Isabella. "Why Don't Cubans See Any Improvement? Skepticism Grows over New Solar Parks." *CubaHeadlines*, 2 Apr. 2025, www.cubaheadlines.com/articles/300064.
7. Team, CiberCuba Editorial. "The UNE Inaugurates Solar Parks in Cuba, but the Electricity Deficit Remains above 1,500 MW." *CiberCuba*, 9 June 2025, en.cibercuba.com/noticias/2025-06-09-u1-e42839-s27061-nid304640-une-inaugura-parques-solares-cuba-deficit-electrico.
8. "The National Electric Grid and the Future of the Cuban Economy (Revisited) | Cuba Capacity Building Project." *Columbia.edu*, 18 Nov. 2024, horizontecubano.law.columbia.edu/news/national-electric-grid-and-future-cuban-economy-revisited.