Traffic Light Control with Reinforcement Learning

Presenter: Devanshi Malaviya ·Computer Information Systems Department ·Mentor: Dr. Sarbani Banerjee

Abstract

The rapid increase in automobiles in the past few years has led to traffic congestions all over the world. This forces drivers to sit idly in their cars wasting time and fuel. Current traffic light control policies are not optimized which leads to people waiting in their cars for nonexistent traffic and extended travel time than necessary. In the US, on an average, people spend about 100 hours in traffic congestions per year. The current research project will focus on reinforcement learning to optimize the traffic flow to reduce the travel time of drivers. It can be done by building an environment where every intersection has knowledge about the number of the vehicles and their speed as they approach the intersection. Simulation of Urban Mobility or SUMO will be used to build a traffic simulator. Reinforcement Learning works on state and action policies which will allow the traffic lights to make optimized decision based on their current state. It will balance the exploration and exploitation to make sure that the model is not overfitting, and every lane is given importance according to how busy it is. For every state, it gets a reward if it reduces the travel time, and the goal of the model is to collect as many rewards as it can. Therefore, at the end this project will attempt to get the most optimized simulation. The Python packages used in this project are Keras, Tensorflow and NumPy.

Packages

SUMO - http://www.eclipse.org/sumo/
TensorFlow - https://www.tensorflow.org/
Keras - https://keras.io/
NumPy - https://numpy.org/

Findings

Figure 1 - Cumulative Delay vs Episode

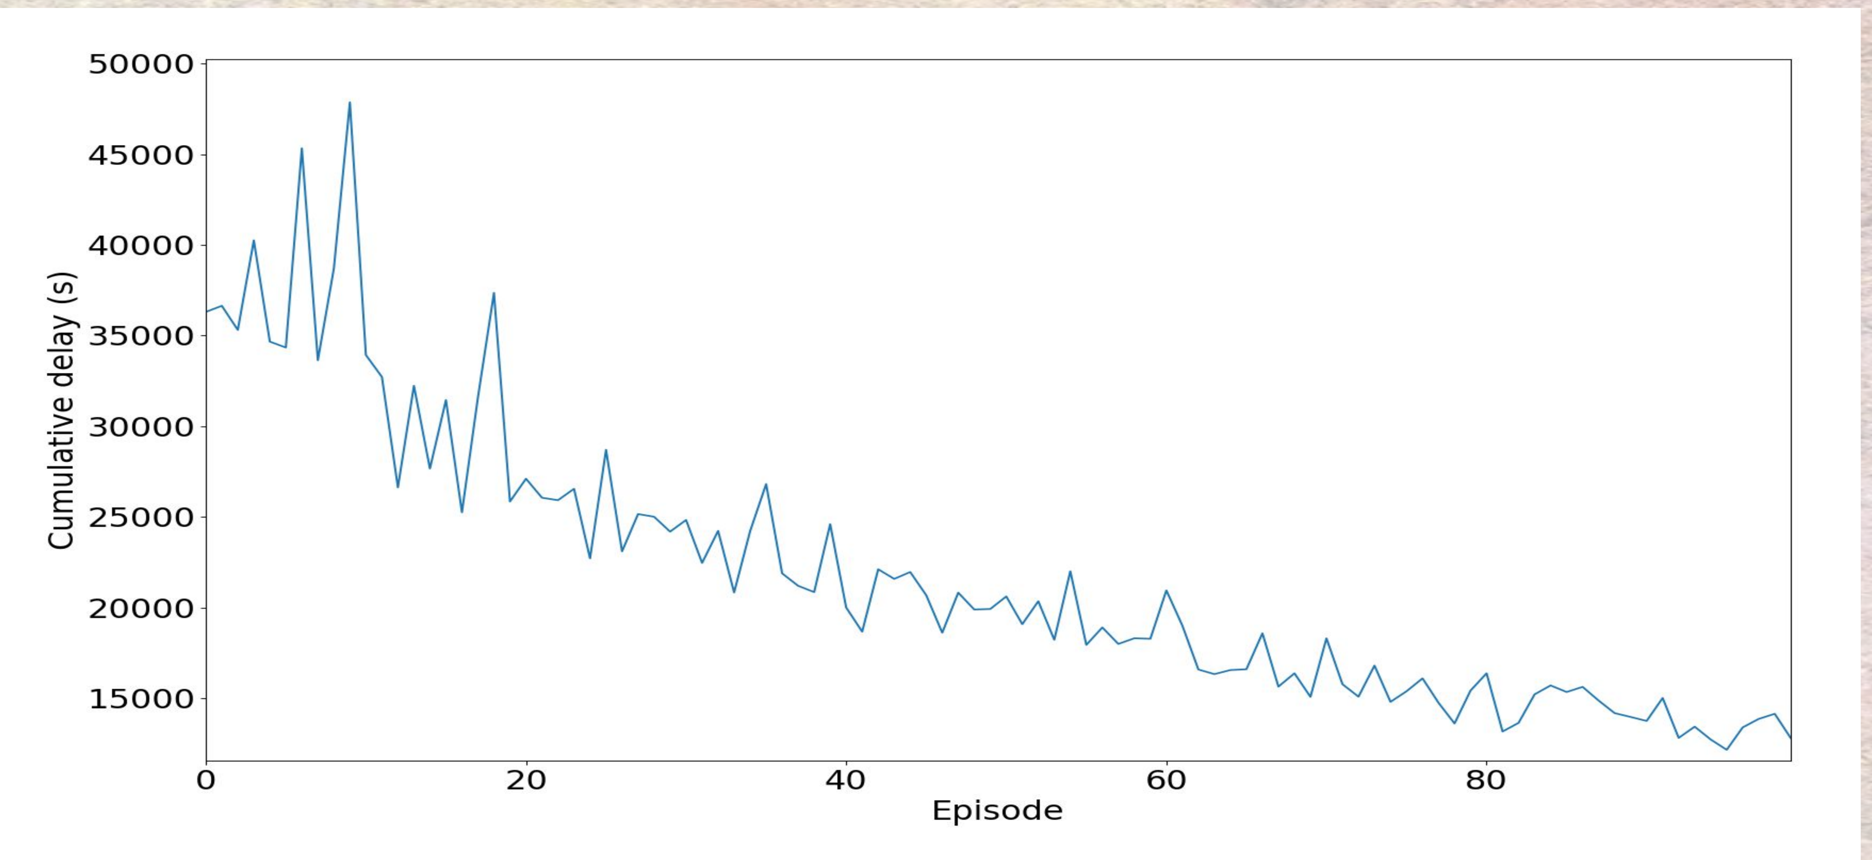

Figure 2 - Queue Length vs Episode

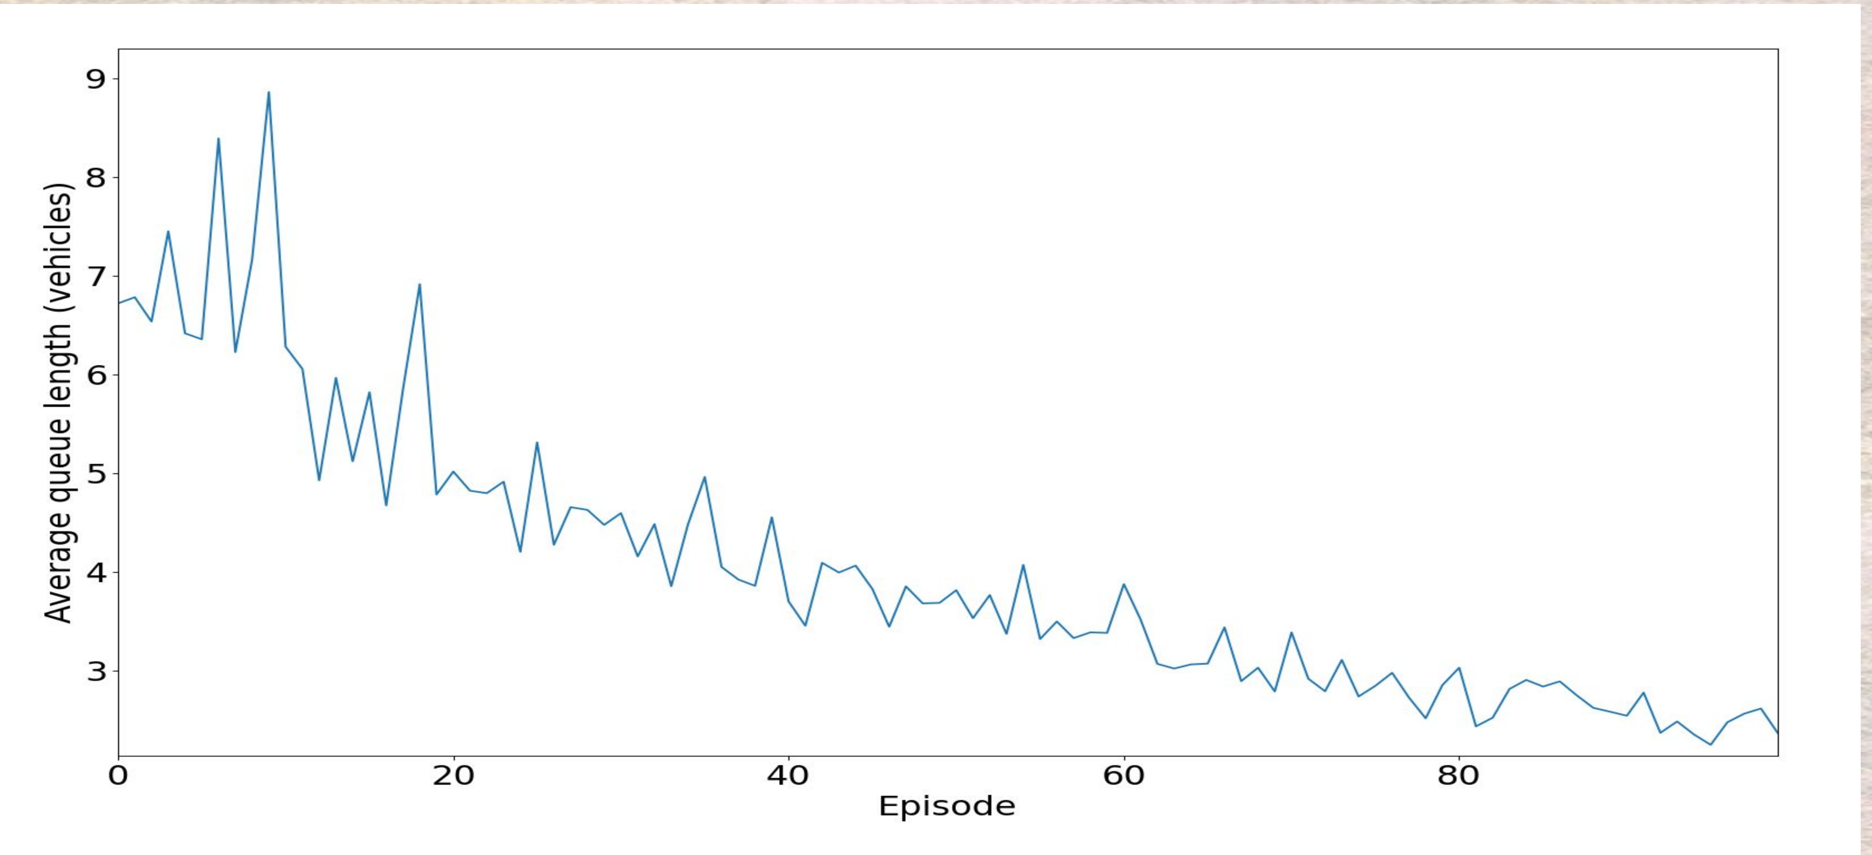

Figure 3 – Simulation of 4-way Intersection

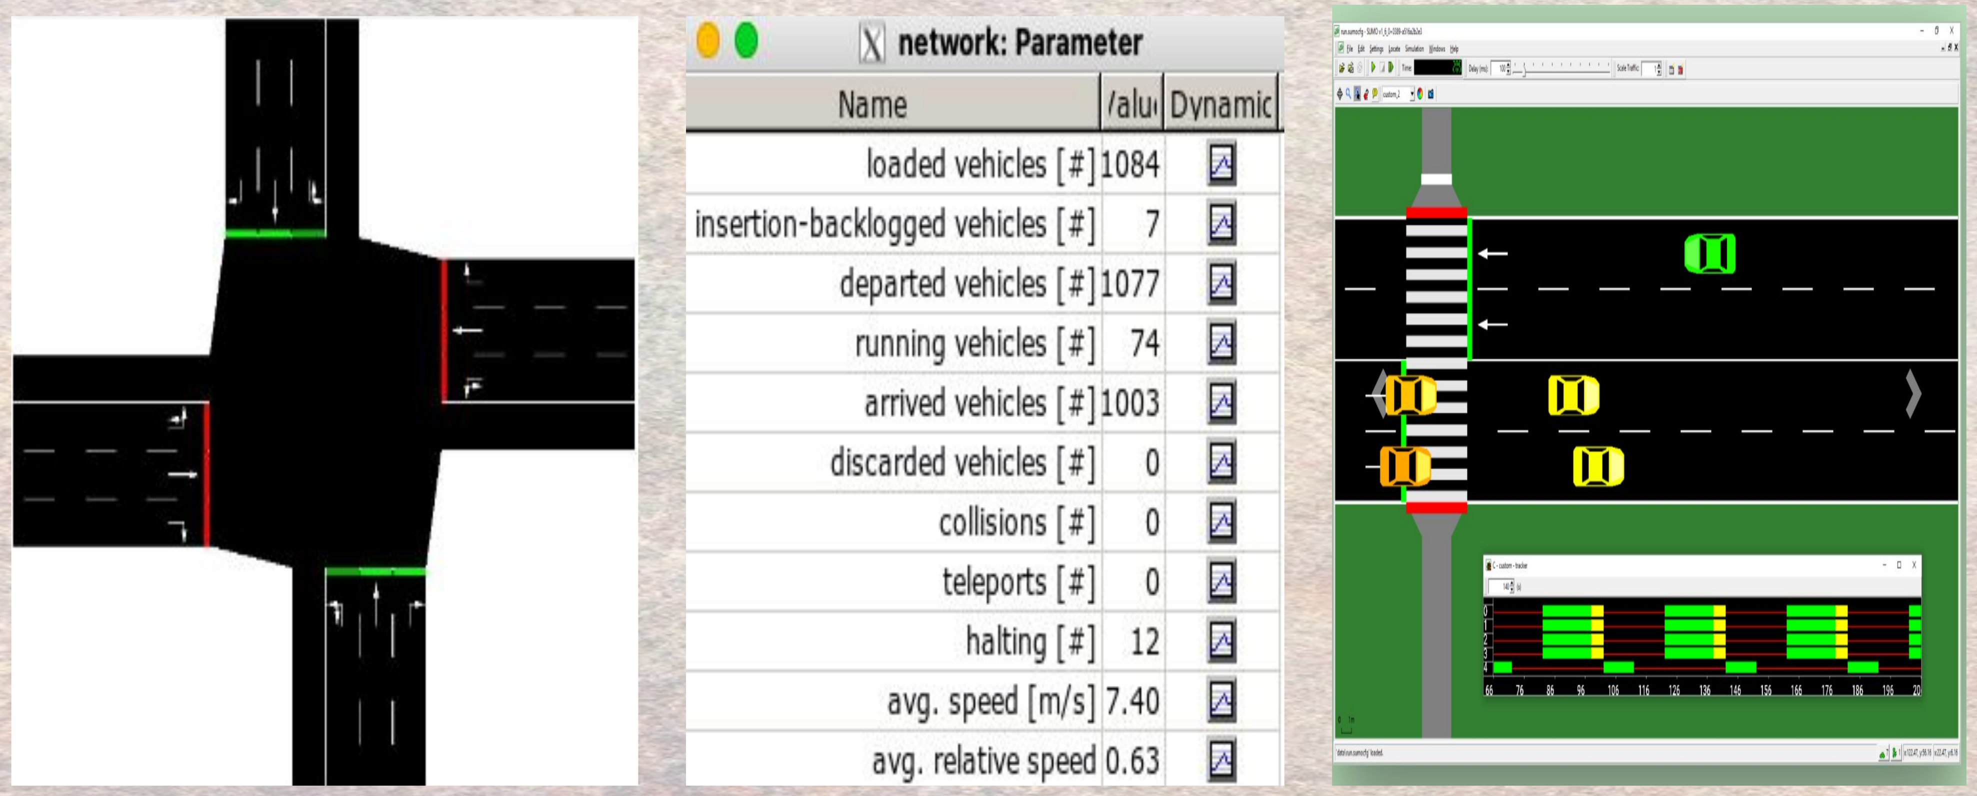

Methods

The project is divided into 2 parts - building the simulation and then running the learning algorithm. In the simulation, a four-way intersection with 4 lanes is selected (Figure 3). The traffic arrival is probabilistic, and it generates the East-West traffic a bit higher than North-South traffic. In every episode of the simulation, there are about 1200-1500 cars, and the simulation runs for 2 hours representing rush hour. Then we run the Reinforcement Learning algorithm that can control the lights of all the lanes. It makes predictions based on velocity of the car, staying time, line length etc. At every state it takes an action and based on that action it get a reward. The program's goal is to collect maximum reward.

Conclusion

The results of this experiment are promising. We saw about 80% decrease in cumulative delays where the cumulative seconds dropped from 47000 secs to just 8000 secs. and 75% decrease in queue length from an average of 8 cars to 2 cars (shown in Figure 1 and Figure 2).

The future research would entail comparing this model with a static system and expand it to multiple intersections depicting a small town or a city.

Acknowledgements

This project is supported by the Office of Undergraduate Research at SUNY Buffalo State and the CIS Department.