NYPD Dispatch Models

Jiaqi Zhang (jz3284), Qiyue Zhu (qz2408), Lucia Nuovo (lvn2105),
Hsu-Sheng Ko (hk3176), Edward Milhem (emm2338)

Team 6

Introduction

- Interested in the tradeoff between the number of vehicles one precinct has and the time between call for service and arrival on scene (waiting time) under different dispatch models.
    - Real World Metrics:
        - ~700s waiting time
        - ~6.5 critical vehicles per precinct (77pcts total in NYC)

- Analyzed the performance of the system by varying:
    - Number of vehicles
    - Incident arrival rates
    - Timeframes
    - Back-up/working policies

End-to-End Response Time

Roll over the graphs to see times for individual segments. All times are displayed as fractional minutes (for example, 1 minute and 30 seconds is displayed as 1.5). Individual segment times are also available in the End-to-End Detail report.

NYC Fleet Daily Service Report: Critical Fleets Summary

Agency	Critical Fleet	Fleet Roster	Target Daily in Service	Actual in Service	Average Over FY 18 to Date
NYPD	Traffic	541	487	494	526

Data resources: https://www1.nyc.gov/site/911reporting/reports/end-to-end-repsonse-time.page
https://www1.nyc.gov/site/operations/performance/fleet-report.page

Data

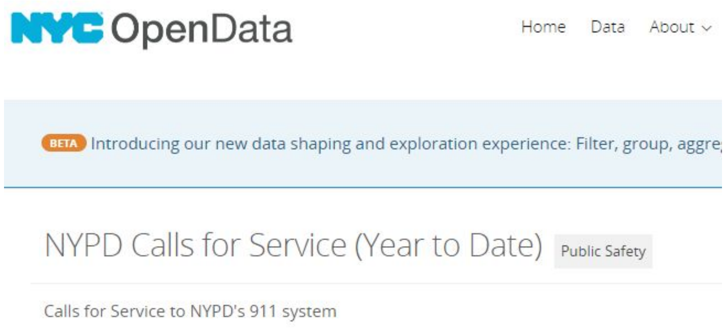

- Incident, Dispatch, Arrival, Closing Times
  - Process = Dispatch - Incident
  - Travel = Arrival - Dispatch
  - Time on Scene = Closing - Arrival
- Incident Locations (Latitude, Longitude)
- Response precinct
  - Focused on Precincts 24, 26, 28
  - 27219, 17855, 27369 incidents respectively

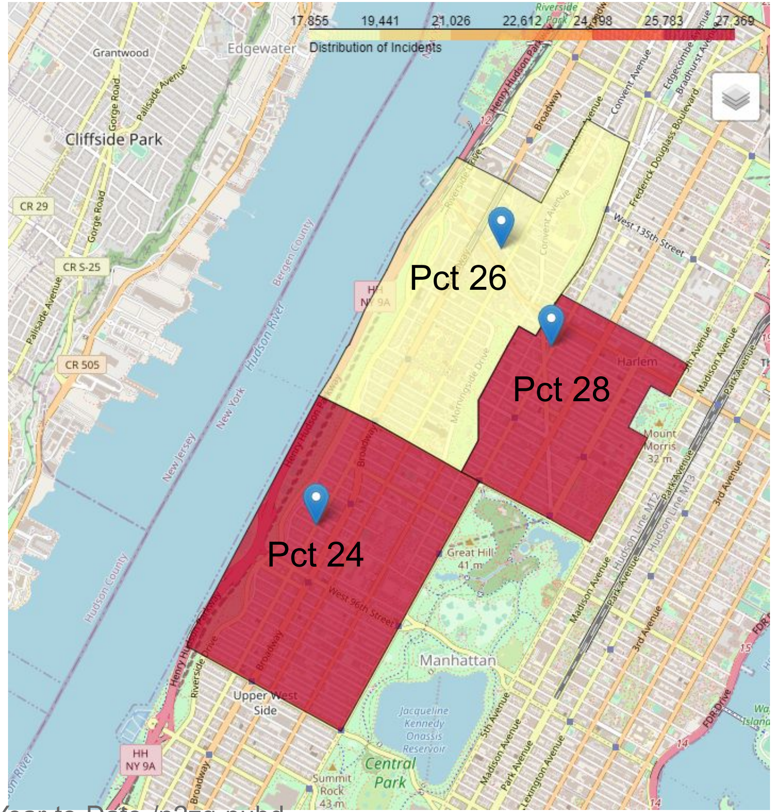

Data resource：https://data.cityofnewyork.us/Public-Safety/NYPD-Calls-for-Service-Year-to-Date-/n2zq-pubd

Assumptions

- Interarrival and service times as poisson processes.
- Instant call processing.
- Haversine distance between start and end locations.
- 40kph travel speed.
- Only 1 vehicle is patched to an incident at once.

Model Implementations

- M/G/n queuing model
- Use *simpy* to simulate the working of police cars and compute the average waiting time of the calls
- Waiting time = Arrival On Scene - Call Added Time
- Simulate 3 dispatch models:
  - 1. No backup dispatch between different stations , return to the station each time
  - 2. No backup dispatch, continuously work for several calls when required
  - 3. Backup dispatch between 3 stations

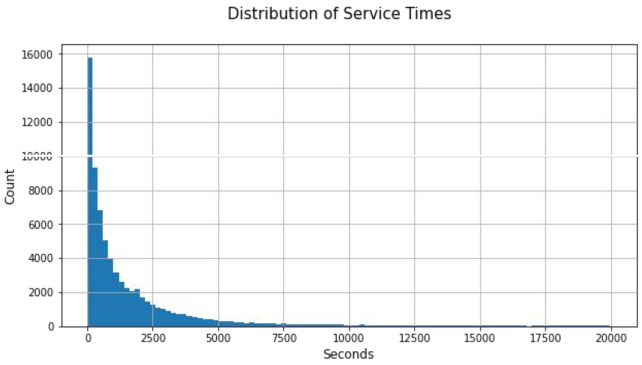

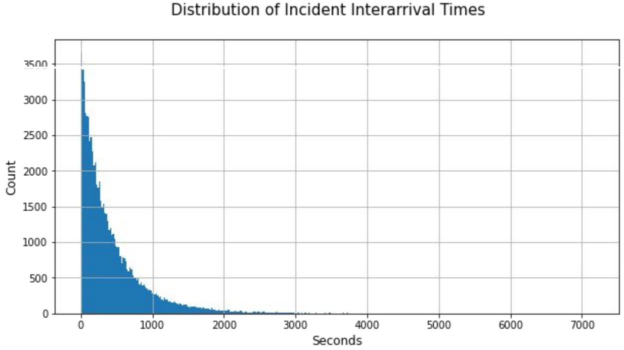

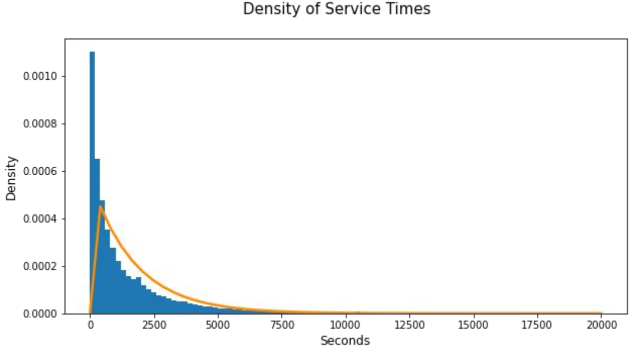

Model 1 : Single precinct, returns to station each time

- Model: the cars in different stations don’t backup each other and the cars have to return to the station each time
- Poor performance at all 3 stations when car numbers = 2
- Increasing car numbers from 2 to 3 can help to make a big improvement
- All 3 stations hit the optimal and satisfy its demands starting from car numbers = 6

Avg waiting time with cars going back			
Car Numbers	Time of 24	Time of 26	Time of 28
2	2156	622	2030
3	369	165	417
4	115	81	136
5	79	73	84
6	71	70	73
7	68	70	70
8	68	71	70
9	68	71	70

Model 1 : Increase and decrease arrival rate

- Increase / decrease the rate parameter by 10% has huge impact when n = 2. They start to converge when n = 4 and hit the optimal at n = 5 to 6 again.

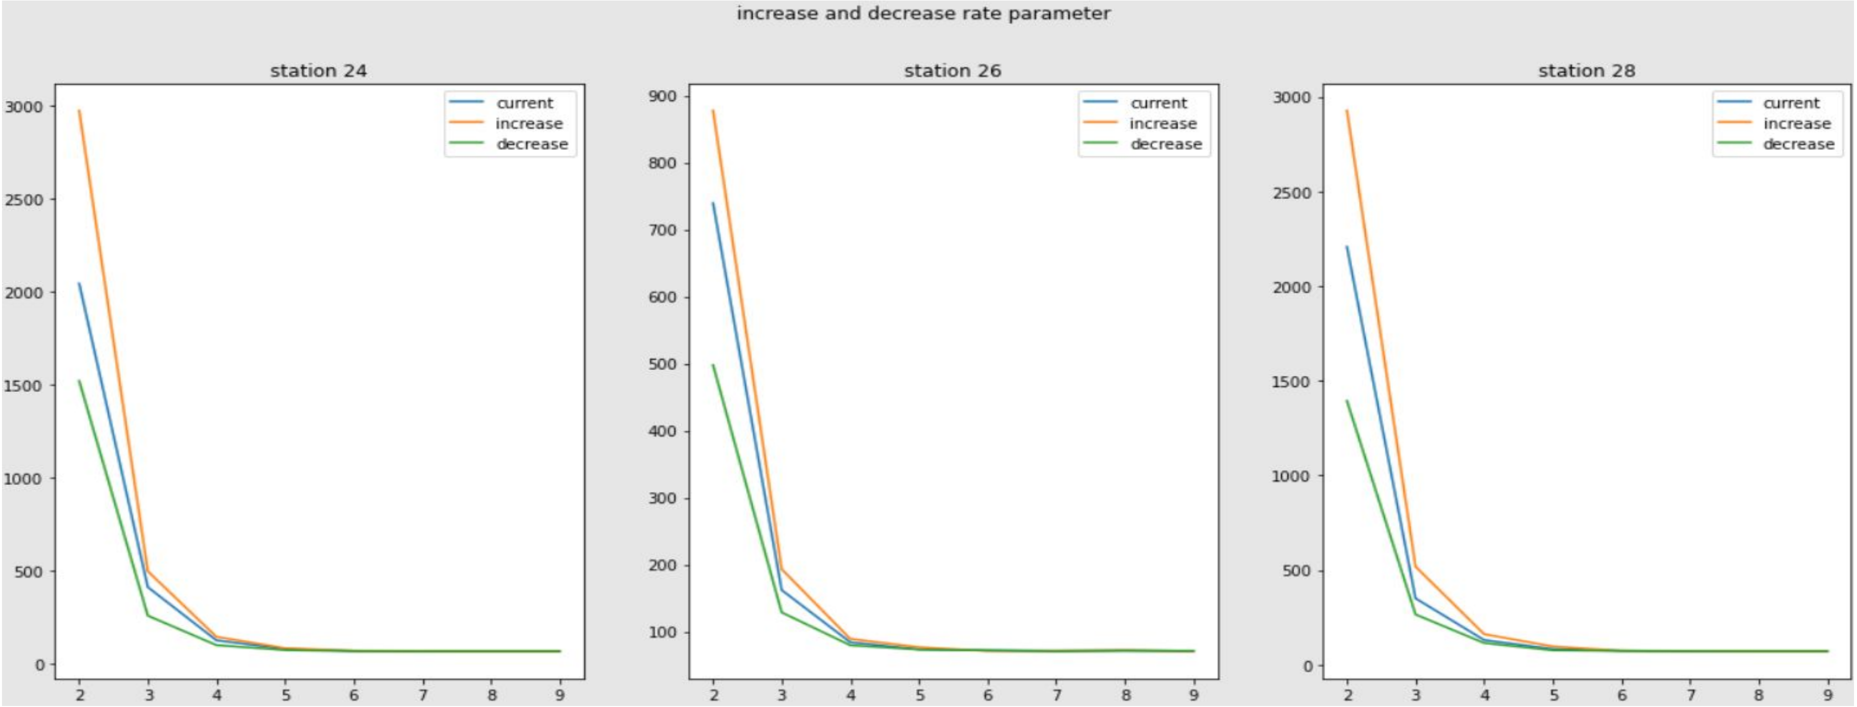

Model 1: Hourly Volatility in Waiting Time

- Incident arrival rate and service rate follow similar daily trends in all three precincts
- Daily service time: (Mean, Variance)
  - Precinct 24: (68.16, **4.41**)
  - **Precinct 26: (71.93, 14.29)**
  - Precinct 28: (79.86,**6.92**)
- Precinct 26 has significantly higher variance than 24 and 28.
- By incorporating a backup system this volatility could potentially be reduced: precinct 26 has the most to gain from a backup system when it comes to volatility reduction

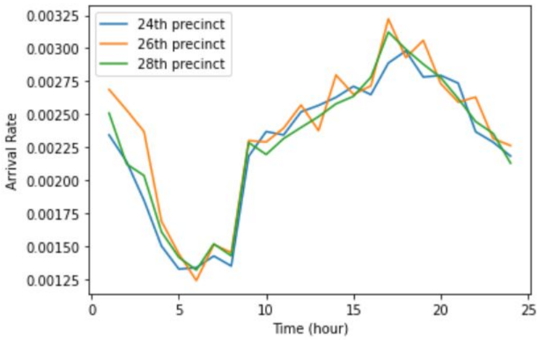

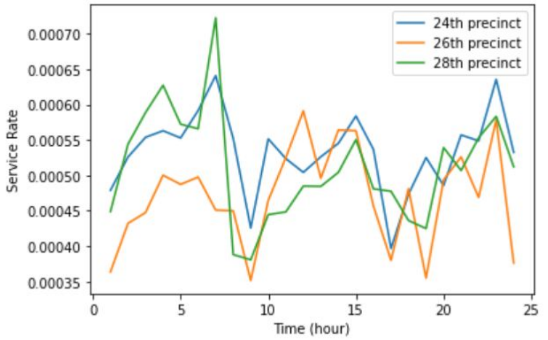

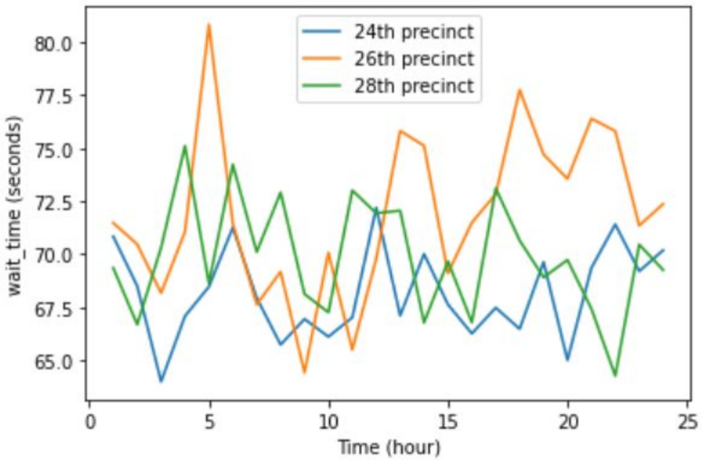

Model 2 : Single precinct but work Continuously

- Model: the cars in different stations don’t backup each other and a single dispatch can run up to 4 consecutive jobs before returns.
- We found that this working policy only has small improvements when n = 2 and 3.
- Once again, it hits optimal for all stations starting n = 6
- We tried working continuously with 2 and 6 jobs as well, only minor differences were found compared with 4 jobs

Avg waiting time with cars working 4 continuous			
Car Numbers	Time of 24	Time of 26	Time of 28
2	2046	603	1864
3	341	149	314
4	125	79	138
5	81	73	82
6	70	71	72
7	69	72	70
8	68	71	70
9	68	70	70

Avg waiting time with cars going back			
Car Numbers	Time of 24	Time of 26	Time of 28
2	2156	622	2030
3	369	165	417
4	115	81	136
5	79	73	84
6	71	70	73
7	68	70	70
8	68	71	70
9	68	71	70

Model 3 : Stations backup each other, returns each time

- Model: cars in different stations can backup each other but the cars have to return to the station each time
- Back-up policy can make huge improvements, system performs much better when n = 2 and 3
- With the new policy, stations could hit the near optimal when n = 4

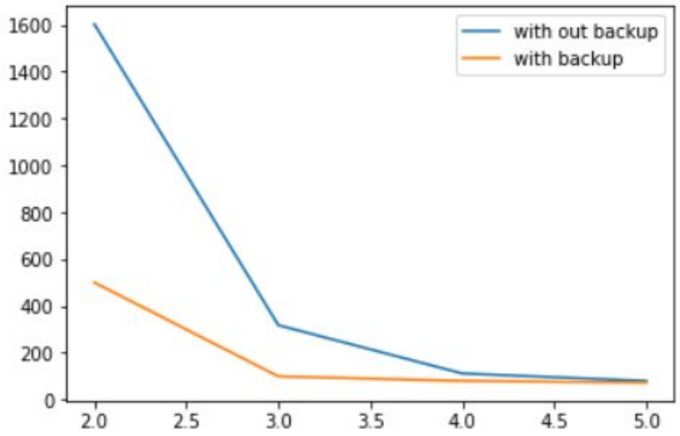

Avg waiting time with/without back up		
Car Numbers	Time with back up	Time without back up
2	499	1603
3	98	317
4	79	111
5	71	79

Practical value

- Help the NYPD to assign proper number of police vehicles to different stations according to the historical reports to increase efficiency.
- Our model can easily be generalized to larger number of stations and larger geographical areas.

Future improvement

- Better estimates of arrival, dispatch, travel, and service times.
  - Time can vary by incident type (e.g. critical vs non-critical).
  - Incorporate real-world travel times using Google Maps API.
- Take the emergency level of the calls into consideration.
- More evaluation metrics like the workload.
- Incorporation of a patrol model allowing vehicles to patrol around in those prone areas to cut the travel time.

Thank You!

Questions