

Optimization of VDOT Safety Service Patrols to Improve VDOT Response to Incidents

A Technical Report submitted to the Department of Systems and Information Engineering

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Julie Gawrylowicz
Spring, 2020

Technical Project Team Members

Elizabeth Cambell
Emma Chamberlayne
Colin Hood
Allison Hudak
Matthew Orlowsky
Emilio Rivero

On my honor as a University Student, I have neither given nor received
unauthorized aid on this assignment as defined by the Honor Guidelines
for Thesis-Related Assignments

Optimization of VDOT Safety Service Patrols to Improve VDOT Response to Incidents

I. Introduction

With the responsibility of building, maintaining, and operating transportation infrastructure across the state, the Virginia Department of Transportation (VDOT) must balance the priorities of safety and economic value with each new initiative brought to the state. Noticing a need for roadside assistance and traffic management, VDOT deployed a fleet of Safety Service Patrols (SSPs) in the 1960s to patrol interstates for drivers in peril and help clear incidents quickly. SSP operators are routed to patrol areas of the greatest need. They continuously patrol between set mile markers during specified hours of the day and stop to assist citizens or they are dispatched to a scene. These mile marker and hour designations are known as patrol schedules. Typically, SSPs stay within their assigned patrol schedule, but they are sometimes permitted to travel off their route if they detect or are notified about problematic incidents close by. In addition to patrol, police officers and other state personnel can request SSP assistance. The use of SSP operators has helped reduce congestion while minimizing secondary crashes [1]. Since their introduction, VDOT has expanded SSP coverage to include 846 miles on all major interstates in Virginia within the Central, Eastern, Northern, Northwestern, and Southwestern regions of the state [2].

VDOT aims to ensure that their current patrol schedules optimize their limited resources while also maximizing their impact. The resources include a fleet of 156 vehicles and a \$2.9 million budget [3]. Current patrol routes are county based and do not employ the use of any optimization or analytical algorithm, leading to potential inefficiencies in time and resources. With the goals of both VDOT and SSPs in mind, we are seeking an optimal schedule to distribute the covered mileage on Interstate-95 into routes that minimize the time vehicles are stranded in need of assistance. The I-95 roadway includes current patrol routes in the Northern and Central regions

of the SSP program. VDOT has accrued vast amounts of incident and traffic data that make a more analytical selection of route schedules within reach. Using complex algorithms, optimization of their current route schedules could lead to faster response times, thus increasing safety and reducing costs.

II. Data and Evaluation Metrics

A. Data

The analysis is based on two main datasets: incident data and traffic volume data. The incident data was created by combining four datasets that were provided by VDOT. Data pertaining to roadway incidents was reported by the existing SSP program. Incidents used in analysis were recorded from August 2017 through October 2019 along I-95, I-195, I-295, I-395, and I-495. The current SSP program includes 14 routes along I-95 and the connecting interstates, with 9 focused

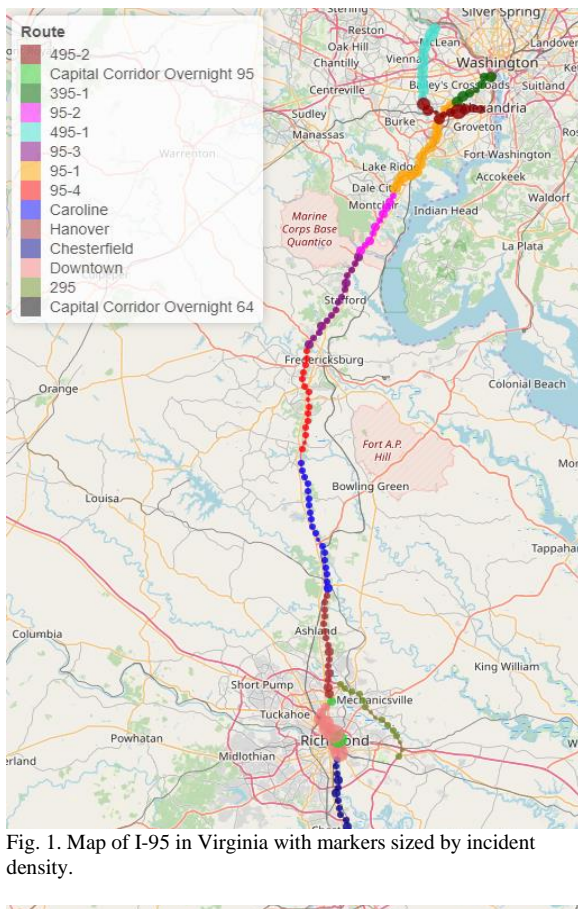


Fig. 1. Map of I-95 in Virginia with markers sized by incident density.

exclusively on I-95. Recorded incidents that had the same mile marker, date, and time were deemed duplicate entries and thus removed from the dataset. After this, we had a working dataset of 88,703 records with 47 attributes that were identified as useful for further analysis. Key attributes for this analysis include: date and time of incident, mile marker, route, direction, and SSP response time. The interstate crosses through two of the state's largest metropolitan areas: Richmond and Northern Virginia. Naturally, these urban areas see the greatest

spikes in traffic density and incident records. The mile markers of I-95 span from 0 to 177 crossing several counties throughout the state. Figure 1 shows the volume of incidents along SSP patrol schedules. The size of the markers show the volume of incidents occurring at that mile marker. It can be seen that in the Northern Virginia and Richmond areas, more incidents have occurred.

To understand typical traffic patterns along I-95, we used traffic volume data. VDOT's Traffic Counts program collects information with regard to hourly traffic density monitored by mile markers in Virginia. Using the hourly volumes, estimates were calculated for the average daily traffic at each mile bin for every day of the week. Estimates were obtained from 2014 to 2018 for I-95, I-195, I-295, I-395, and I-495.

Overall, the data describes SSP activity and incident frequency throughout the state. Over half the incidents SSPs respond to are disabled vehicles, with the second most common incident type being a vehicle accident.

B. Metric of Evaluation

VDOT has a wide array of information surrounding traffic incidents, SSP operations, and daily traffic volumes. However, currently there is not a single, comprehensive metric to evaluate SSP performance in response to incidents. Therefore, we developed a new metric that fuses these information sources in order to capture the overall influence of the program. The SSP program's goal is to quickly respond to incidents in order to minimize the amount of time a vehicle is disrupting traffic. The ideal scenario for the SSP program would be to instantly detect every traffic incident. As this is not always possible, patrol schedules should be designed such that SSP patrollers can minimize the amount of time that drivers in need are waiting for assistance. Thus, the metric for route optimization is the total time a vehicle is waiting for SSP assistance or total time stranded. The metric combines the response time to incidents and the amount of incidents by

summing the estimated incident counts per mile marker multiplied by their corresponding response time. This motivates the creation of shorter routes over mile markers with high incident counts to achieve faster response times. Although SSPs cannot control incident rate, they can control how long they take on average to respond to an incident by having shorter routes. Overall, this metric gives a comprehensive score to route schedules which allows for comparison of route schedules and optimization.

III. Methodology

Our goal is to leverage the total time waiting metric in order to develop optimized patrol schedules along I-95. To do this, we created an optimization model using a genetic algorithm (GA). Inspiration for the use of a GA comes from police patrol routes which have used GAs in the past to optimize police response to crime [4]. Police routes aim to fulfill a similar goal to SSPs since “police must proactively patrol and prevent offenders from committing crimes but must also reactively respond to real-time incidents” [4]. Similarly, SSPs patrol for undetected incidents and respond to incidents where they are dispatched. In addition, a GA was chosen as it is flexible to fit with the scheduling constraints for the routes and their evaluation. The flexibility allowed for combining and changing routes to score a variety of options while maintaining the constraints. GAs do not guarantee that a global minimum is found, but they can provide an improvement. Any improvements can be valuable to VDOT as other optimization methods have not been tried for their route schedules. For example, the improvements could be used as a starting point for further optimization efforts or provide general insight to factors influencing route performance.

A GA is based on the phrase “survival of the fittest” often used in evolution [5]. It uses the notion that an offspring's set of optimization rules can derive better results than its parent by pulling from solutions of the parent's set of optimization rules [5]. The fittest offspring in this case, which

are the routes with the lowest wait time, get passed to the next generation and influence the search towards an optimal solution. GAs are useful in a scenario when there is limited information about certain separate parts of the solution [5]. For detecting incidents, VDOT has reliable information on daily traffic volumes and past detected incidents, but lacks information on incidents in the future which makes finding a solution more difficult.

To calculate the total time stranded metric, we created models to estimate incident rates and response time of an SSP to an incident by mile marker. The incident and response time estimates serve as input to the GA and are the key aspects of scoring the route schedules.

A. [Estimating Incident Rate](#)

VDOT collects information about all the incidents that SSP patrollers respond to across the state, however the transportation agency may not be made aware of every incident that occurs. We determined that regions outside of SSP coverage had disproportionately lower incident counts compared to areas of similar traffic volume that are currently covered by SSP patrol schedules. This bias in the incident reporting illustrates the effectiveness of SSP patrollers when they are deployed to new coverage areas, but to find the optimal locations for patrol we must use statistical modeling to address these concerns and estimate incidents in areas that previously have not been covered.

In order to predict incidents, we used VDOT daily traffic data. VDOT collects data on the average daily traffic volume broken down by the time of day and mile marker. We hypothesized a relationship between the daily traffic and the number of incidents at a specific mile marker and hour of the day. Below is a scatterplot displaying the relationship where each point on the graph represents a mile marker at a specific hour.




Fig. 2. Scatterplot of incidents vs. traffic volume.

The scatter plot indicates two trends in the data. We previously noticed a difference in the data collection processes of the Northern and Central region, so we proposed those differences may describe the two trends in the data. Below is the scatterplot displaying the same relationships but controlling for regional differences.




Fig. 3. Scatterplot of incidents vs. traffic volume controlling for region.

These trends show a relationship between traffic volume and incidents in both regions. However, the trend for the Northern region suggests that this area experiences fewer incidents than the Central region. Based on discussion with VDOT and the strong positive trend in the Central region, the low incident rate in the Northern region is unexpected. The daily traffic volumes that are experienced in Northern Virginia do not rationalize the vastly lower levels of observed

incidents as compared to the Central region. In fact, more incidents would be expected in the Northern region based on the traffic volume. The discrepancy between the two regional models for incidents may be explained by differences in data reporting. Without reason to believe the Northern region has less incidents, we moved forward with predicting incidents using only the Central region. We believe that the reporting in this region is closer to the true level of observed incidents.

We determined that the most appropriate model to predict incident counts based on average daily traffic volume data is a Poisson regression model. Poisson regression is designed to predict count data, like incidents. Due to this, it is the most commonly used approach in traffic modeling. Basu and Saha described useful modeling techniques for traffic crashes that provided a basis for our investigation [6]. The ultimate model selected is:

$$\log(\hat{u}_i) = 1.603 + 0.0000351W_i \quad \square\square\square$$

Where u_i represents the predicted number of incidents and W_i represents the reported average daily traffic volume.

B. Expected Response Time

The VDOT reported incident data included a field for response time. These response times were measured from the time the SSP knew the incident location to when the SSP arrived at the scene. This resulted in reports of zero and one minute response times which indicated that the SSP detected the incident and therefore was already on scene when they knew the incident location. This does not represent the metric of time a vehicle is waiting for assistance before SSP arrival. This motivated the following model for expected response time.

The expected response time for an incident occurring within route coverage is modeled based on route length. The route length is the number of mile markers covered multiplied by two since

the SSPs travel the route in both directions. At each mile marker on a given route, the time in minutes required to travel to any other mile marker can be calculated by multiplying the number of miles between the two mile markers by minutes the SSP can travel per mile. We set the speed of an SSP to a constant 55 miles per hour, since it is the minimum speed limit throughout I-95. Then, to calculate minutes traveled per mile, we divided 60 by 55 which resulted in about 1.09 minutes needed to travel one mile. At the time of an incident, we assume that the probability of an SSP being at any given mile marker is equal. This simplifies the response time to an average of travel times along a route. The minimum distance needed to travel is one mile and the maximum is the entire route length. The minimum and maximum times correspond to an incident happening just before or after an SSP passes it. Given route length n , the expected on route response time can be written as:

$$\frac{60}{55} * \frac{1}{n} \sum_{i=0}^n i \quad \square \square \square$$

The expected response time for an incident occurring outside of route coverage was calculated as a linearly interpolated value between the modelled response times of the closest routes in either direction plus an additional dispatch time. The dispatch time accounts for the time taken for an alternative detection source to discover the incident and report it to an SSP. The interpolated value then accounts for the response time of the dispatched SSP from the closest route to the incident outside of coverage.

The response time is a maximum expected time as it only considers SSP detections for on route coverage. There are other methods of detection for on route incidents, such as the Virginia State Police, which could notify an SSP dispatcher and expedite the SSPs arrival. Another assumption is that the SSP will not see an incident in the other direction of travel and be able to turn around early.

C. Optimization

In order to schedule new optimal routes, we employed the use of a GA. GAs are designed to mimic Darwin's theory of evolution through natural selection. The process designs an algorithm where "the fittest individuals are selected for reproduction in order to produce offspring of the next generation" [7]. The GA has five components: an initial population, fitness function, selection, crossover, and mutation. The initial population corresponds to randomly generated route schedules on the exit numbers along I-95. Nine routes were used to match the current amount on I-95 and they needed to be at exits to ensure that the SSP vehicle can efficiently loop around the route. The number of start and end point groupings represents how many routes are in each schedule. The translation of the vector input (v) by mile marker used in the GA to the routes (r) in a schedule is in Equation 3.

$$v = [51, 73, 73, 83, \dots, 170, 177] \quad (3)$$

$$r_1 = \text{I-95 Exit 51-73}$$

$$r_2 = \text{I-95 Exit 73-83}$$

$$r_n = \text{I-95 Exit 170-177}$$

The overall set of random schedules should be comprehensive to ensure a varied initial population. We generated a total of 970 schedules in the initial population. The schedules were created by randomly selecting the start and end points from a list of all of the exit numbers or subsets of the exit numbers to increase the variety of coverage areas.

Next, the fitness function evaluates the metric of performance, defined as total time waiting for SSP assistance. The fitness function calculates the response time per mile marker based on the model described previously. It then sums the multiplication of the estimated response time and the corresponding number of incidents per mile marker to obtain the score. After each schedule is

scored, selection occurs by picking the top half of the population as the parents of the generation. Since our initial population is 970, the algorithm will select the top 485 route schedules with the lowest scores.

Crossover creates the second half of the generation, called offspring, by combining information from the parents. This is done by merging the first half of the routes from one parent and the second half of the routes from another parent. If the schedule assumptions are not preserved in an offspring, such as no overlapping routes, the schedule is modified. In the case of an overlap, the schedule will have less routes than required. This is solved by splitting the longest routes until the necessary number of routes is reached.

Finally, mutation occurs on the offspring to preserve the previous best solutions, the parents. Mutation ensures diverse solutions and explores other potential route options that may not have otherwise been considered by the crossover function. To mutate the offspring, we change one of the start or end exit numbers of a route in each schedule at random. The exit number can change to the next or previous one or two exit numbers. Again, we check to ensure that the schedule assumptions are preserved after the mutation. Since parents are not mutated, the next generation will either find a more optimal solution through crossover and mutation, or maintain the same most optimal solution from the parents of the previous generation. This ensures that the score will not become worse over the generations. If crossover does not find a better solution, mutations on the parents would risk the previous best score to increase and thus make the next generation worse than before.

The GA continues to run until the termination condition is reached. This occurs when the score of the most optimal schedule remains consistent for 50 generations. The final outcome is a route

schedule based on exit numbers that minimize the total estimated time a vehicle is stranded without SSP assistance.

IV. Results

The majority of VDOT's routes on I-95 run 24 hours or 16 hours a day and seven days a week with eight hour shifts. This does not account for potential variance of traffic patterns on the weekend or weekday or during different hours of the day. The results aim to determine whether there can be a significantly improved overall schedule and if specific schedules based on traffic patterns should be considered.

Each schedule was created using a subset of estimated incidents that represent weekday, weekend, and each day of the week and are further divided into three eight hour shifts. This totals 27 subsets of estimated incidents per mile marker to input into the GA. The overall schedule was built using the sum of the estimated incidents.

The significance between the fitness scores of the initial and overall optimized schedule was evaluated by a one sided Wilcoxon signed-rank test. This tested if the median of the differences between the current and optimal schedule scores can be considered greater than zero or not. A median greater than zero would indicate the optimal schedule has a lower, and therefore better, score of total time waited. The sample consisted of the scores of the two schedules over each of the 27 estimated incident subsets. The assumptions of this test include paired samples, independent differences, and symmetric distribution of the paired distribution. The results showed that a greater difference was significant at the 0.01 level. This shows that the overall optimized schedule can increase performance despite differences in estimated incidents from varied levels of traffic. The average percent difference between the scores was a decrease of 13%. Based on the metric, this

corresponds to saving an estimated 442,000 minutes of total time waiting for SSP services over one year.

The routes and average route lengths of the specific schedules provide insight on how schedules could be adjusted to improve performance depending on the time of day. The overall route and day to day specific routes cover a similar region to the current coverage over mile markers 50 to 177. The average route length of these schedules ranges from 20 to 25 miles, which is lower than the current average of 28 miles. The greatest difference was in the routes generated to optimize response to the subset of weekday incidents. These schedules were concentrated from mile markers 70 to 177 with an average route length of 10 to 15 miles. This is reasonable as this goes from the beginning of Richmond and up through Northern Virginia, which are areas of higher traffic.

Overall, these results recommend further consideration of maintaining separate weekend and weekday route schedules due to high traffic volume. It also supports VDOT's current practice of keeping routes less than 30 miles and not having routes on the lower part of I-95.

V. Conclusion

Although the results can potentially save over 400,000 minutes of drivers in need of help, the GA can reach far beyond the drivers of Virginia. The data-driven solution VDOT leaders are seeking can help save lives, improve safety, reduce congestion, and reduce fuel consumption. The GA can improve services and lower inefficiencies for any department of transportation looking to optimize their efforts. The system that was built can be applied to other patrol services like ambulances, food trucks, and police officers. With slight modifications, the system outlines the process of implementing a GA, from the data cleaning, boundary definitions, response time estimation, and incident predictions.

VI. Limitations

A large portion of time was spent cleaning the data and ensuring that it was fit to apply to a model. There is a difference in reporting methodology between the regions associated with I-95 routes that created a disparity in data which needed to be resolved with statistical modeling. This process, though fit to Interstate 95, has not been tested on other roadways. This creates an individualized use case that has not yet been evaluated on other interstates. Additionally, GAs create solutions that are not guaranteed most optimal, but rather work towards an optimum. Therefore, it is not possible to say that the results are the best possible solution, but rather an improvement in minimizing total time waiting.

VII. Future Work

This model for the current route can be expanded upon in the future to be more robust and applied to more routes. The model in place should fit to new data in other regions and expand to other routes and interstates. With more accurate reporting of data and equal quality reporting of data amongst all regions, the model can be expanded and iterated over to create a more precise result. The use of other tools that the VDOT records can be supplemental to this work. Closed Circuit Televisions (CCTVs), or traffic cameras, are placed along much of the interstate and can be used as supplementary information to provide real time information to SSPs. In the future, if these improvements are implemented, the developed models for predicting incidents and response time can be tested against new data to evaluate their performance. If the models are not accurate to the region of patrol, they will need to be adjusted accordingly based on this more precise data.

Acknowledgment

The authors would like to thank the Virginia Department of Transportation for sponsoring this research and in particular Katie McCann and Daniel Truitt for their time and support.

References

- [1] Cetin, M., Khattak, A., Wang, X., et al. (2011). *Primary and Secondary Incident Management: Predicting Durations in Real Time* (p. 83). Civil and Environmental Engineering Department Old Dominion University.
http://www.virginiadot.org/vtrc/main/online_reports/pdf/11-r11.pdf
- [2] *Safety Service Patrol—Travel* / Virginia Department of Transportation. (n.d.). Retrieved April 14, 2020, from <https://www.virginiadot.org/travel/safetypatrol.asp>
- [3] Mitzel, C. (n.d.). *Here's how one VDOT program helps drivers stuck on the interstate*. The News Leader. Retrieved April 14, 2020, from <https://www.newsleader.com/story/news/2019/07/02/vdot-safety-service-patrol-expands/1621582001/>
- [4] Dewinter, M., Vandeviver, C., T., Witlox, F. (2020, March 9). Analysing the Police Patrol Routing Problem: A Review. *International Journal of Geo-Information*
- [5] Ross, P. (1997) Commentary—What Are Genetic Algorithms Good at?. *INFORMS Journal on Computing* 9(3):260-262. <https://doi.org/10.1287/ijoc.9.3.260>
- [6] Basu, S. & Saha, P. (2017). Regression models of highway traffic crashes: A review of recent research and future research needs. *Procedia Engineering*, 187(2017), 59-66.
<https://doi.org/10.1016/j.proeng.2017.04.350>
- [7] Mallawaarachchi, V. (2017, July 7). *Introduction to Genetic Algorithms—Including Example Code*. Medium. <https://towardsdatascience.com/introduction-to-genetic-algorithms-including-example-code-e396e98d8bf3>