

A COMPARISON OF EVALUATION METHODS FOR POLICE PATROL DISTRICT DESIGNS

Yue Zhang
Samuel H. Huddleston
Donald E. Brown
Gerard P. Learmonth

Department of Systems and Information Engineering
University of Virginia
Charlottesville, VA 22904, USA

ABSTRACT

Police patrol district design presents a multi-objective optimization problem with two goals: minimizing workload variation between patrol districts and minimizing the response time for officers responding to calls for service. We evaluate three different methods for scoring district designs: a closed form probability based approach, a discrete-event simulation based on hypercube models for spatial queuing systems, and an agent-based simulation model. We find that all methods provide similar evaluations when service demand is low enough that cross-boundary support is infrequent. However, when the demand for service routinely exceeds the supply available within districts, only the agent-based simulation model accurately represents the resulting complexities and significantly changes the evaluation scores to reflect the behavior of the system.

1 INTRODUCTION

Nearly every police department creates geographic patrol districts (also called patrol sectors or car beats) as a standard management method to enhance the capabilities of the uniformed patrol force (Hale 1980). Better districting plans lead to lower response times, officer's familiarization with their assigned area, more efficient use of personnel, more equal division of workload, a visible police presence, enhanced officer safety, officer accountability, and balanced police response to calls (Hale 1980). Traditionally, these geographic patrol boundaries are drawn by hand based on a police department's knowledge, experience, and the available police resources (Mitchell 1972; Taylor and Huxley 1989). Most police departments also lack a method for formally evaluating and comparing the performance of competing district plans, instead relying on the judgement and intuition of police planners. However, given the complexities of the police districting plan, it is unlikely that an optimal districting plan will be chosen by chance using this method (Curtin, Hayslett-McCall, and Qiu 2010).

Police patrol district design presents a multi-objective optimization problem with two goals: minimizing workload variation between patrol districts and minimizing the response time for officers responding to calls for service. Fast response to citizen Calls for Service (CFS) improves the chances of arresting offenders, increases the chances of identifying and locating witnesses, provides immediate gathering of physical evidence, provides immediate life-saving aid, enhances the reputation of the police department, and increases citizen satisfaction with police (D'Amico et al. 2002; Hancock and Simpson 2009). Therefore, virtually every police department seeks to minimize their average response time to CFS. Workload variation between districts arises because crime (and other CFS) tends to cluster in "hot-spots" rather than being uniformly distributed in the city. Workload variation between districts is often high, with some officers/districts experiencing much higher CFS volume than others. When small districts are created around very "hot" zones, the remaining districts can be quite large, resulting in slow response times for many citizens. Thus,

minimizing workload and reducing response times are often competing objectives, requiring police to select a comfortable trade-off point between the competing objectives.

In this paper, we compare three different methods for evaluating the performance of patrol district designs in this trade-off space: a closed form probability based approach, a discrete-event simulation based on hypercube models for spatial queuing systems, and an agent-based simulation model. We use the selection of a new patrol districting plan for the City of Charlottesville, Virginia, as a motivating example to compare and contrast the different methods for choosing a districting plan. We find that although all three methods provide similar evaluations of the districting plans when the emergency response system is not stressed, the agent-based simulation model more accurately represents the system dynamics when the system is highly stressed and also yields important insights into the system dynamics that the other two methods do not provide.

2 POLICE PATROL DISTRICT DESIGN

Police patrol district design is one of the most important resource allocation problems police departments must consider. As previously noted, this problem is usually solved based on officer experience, judgement, and intuition. However, researchers have proposed several different approaches for optimizing police patrol boundaries in a city. In the first proposed approach, Mitchell (1972) proposed using p-Median clustering to minimize the total weighted travel distance to service expected calls. The problem can also be specified as a classic problem in mathematics. The districting problem can be defined as the aggregation of smaller units (that must be specified) into larger geographic units that form districts. This problem reduces to a classic graph-partition problem with the constraints of contiguity and compactness and has been shown to be NP-hard (Altman 1997; Johnson 1985). Researchers have also proposed heuristic methods for identifying “good” (locally optimal) districting plans including simulated annealing (D’Amico et al. 2002), maximal covering models (Curtin, Hayslett-McCall, and Qiu 2010), genetic algorithms, and stochastic gradient ascent (Zhang and Brown 2013).

Unfortunately, all of the methods specified above require solving optimization problems, some of which are NP-hard. These approaches are well beyond the capabilities of the vast majority of police departments. Instead, we propose that a more reasonable approach would be to provide police departments with the ability to evaluate the performance of districting plans they produce themselves and/or automated methods to generate a large number of possible districting plans. This approach would allow police departments the ability to find the optimal plans within any defined set. While these plans might not be globally optimal, they are likely to provide good performance, especially when compared to plans drawn by hand that consider the resulting performance trade-offs only in the minds of the planners. In this paper, we use an automated method to generate a set of districting plans for a case study and compare methods for evaluating the utility of these plans for department use.

3 CASE STUDY DATA

The Charlottesville Police Department provided the data used as the case study for this analysis. The City of Charlottesville is a mid-size city centrally located in the state of Virginia, USA. The city has a diameter of about 7 miles and a year-round population of about 40,000, which swells to about 66,000 during the academic year due to the presence of a major university. The current districting plan used by the CPD is about 20 years old. The city uses eight city patrol districts, with one car routinely assigned to each patrol district during each patrol shift. The police department operates three shifts a day: morning, evening, and overnight. As is the policy in many police departments, the CPD always dispatches the nearest available car to the scene of a CFS in an effort to minimize response time, rather than relying on each police car to respond to all calls within its’ district.

Figure 1 provides an illustration of how the demand for police assets varies over the 24 hour period. This graph references 330,000 CFS incidents observed over a four year period. During the night shift, the

inter-arrival time for Calls-For-Service (CFS) is high, meaning that the CFS intensity is low. During the day and evening hours, CFS intensity is high, placing greater demands on the police patrols. As Figure 1 illustrates, the response time for police responding to CFS is highly correlated with traffic volume in the city. At night, police can respond relatively quickly to CFS because there is little traffic. During the morning and evening rush hour periods, it takes much longer for police to navigate traffic to the scene of calls for service. The time on the scene for a CFS remains relatively stable over the 24 hour period. The dashed lines in Figure 1 correspond to the three modeling scenarios used to study the performance of the police patrol district designs:

1. Low-Intensity Demand: 5 AM (Night Shift)
2. Medium-Intensity Demand: 7 PM (Evening Shift)
3. High-Intensity Demand: 9 AM (Day Shift)

The low-intensity scenario represents when the system is most idle. The medium-intensity demand scenario represents a time period when all system parameters are near their average values. The high-intensity demand scenario represents the time period when the system is most stressed - during the morning rush hour. Examining these scenarios allows us to evaluate patrol district designs under minimum, maximum, and average conditions.

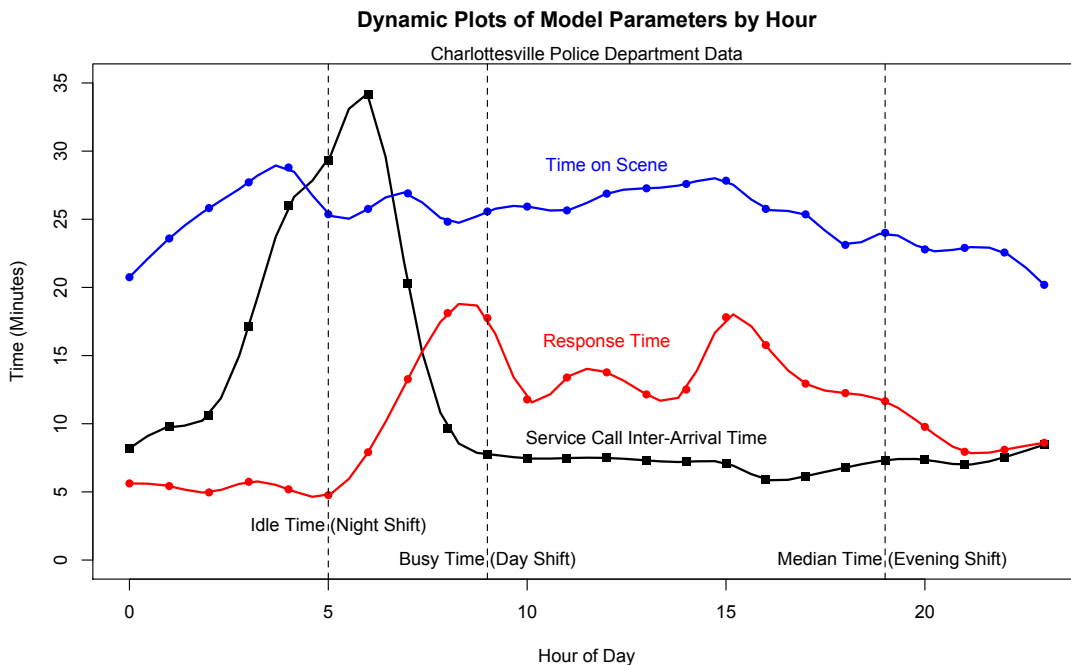


Figure 1: A plot showing how the service call inter-arrival time, response time, and time on service model parameters vary over time in the Charlottesville Police Department (CPD) data set. The dashed lines represent time periods selected for study in the simulation models: a low intensity (idle) period, a high intensity (busy) period, and a period representing the median situation.

Figure 2 provides one of the considered patrol district designs plotted over a map of the city of Charlottesville. As can be seen in Figure 2, we organized the city into 323 atoms (locations) for assignment. Each of these atoms must be assigned to one of the eight police districts. We generated 150 possible police district designs for consideration by the CPD using the procedure outline by (Zhang and Brown 2013).

This procedure develops districting plans that are both contiguous and compact. By definition, the patrol district must be contiguous so that one patrol car can patrol that area without departing. Police desire compact districts to provide shorter travel times within the patrol district. This procedure converts the NP-hard graph partitioning problem into a much more tractable problem: choosing the best of a defined set of options. The limitation of this approach is that the generated choice set is not guaranteed to contain the optimal solution. Rather, the approach we have taken provides a set of reasonable solutions from which we would like to choose the best available.




Figure 2: Visualization of District Plan 21 under evaluation in the agent-based simulation. The eight districts are color-coded. In-use patrol cars are labeled as stars while patrolling cars are labeled with circles. Note that the green and red patrol cars are responding to incidents out of sector because they were the closest available car to the incident at the time of the CFS.

4 COMPARISON OF POLICE PATROL DISTRICT EVALUATION METHODS

Given a patrol district design, there are several different approaches available for assessing the utility of that design. One approach would be to try the different districting plans by asking police patrols to change their patrol sectors every few weeks and assess how well the various districting plans worked. However, the number of possible districting plans that a police agency could try would be very limited and it would take a long time to test even a limited few competing plans. Efficient methods for evaluating patrol district designs without actually testing them in practice are therefore in high demand.

In this paper, we compare three different methods for scoring district designs: a closed form probability based approach, a discrete-event simulation based on hypercube models for spatial queuing systems, and an agent-based simulation model. The closed form probability-based approach greatly simplifies the modeling problem and requires only geographic data in order to make an estimate. It is very simple, fast, and can be applied by virtually any police analyst with access to a GIS system. The discrete-event model takes longer to develop and makes some simplifying assumptions about the problem. However, once developed, it can evaluate districting plans quickly and it is relatively simple to adapt the model to different cities and scenarios. The agent-based approach takes the most development time, is harder to adapt to different cities and environments, and takes longer to evaluate competing plans. However, the agent-based modeling approach provides the most high-fidelity representation of the system and the most flexible modeling framework. We expound on each the three evaluation methods in greater detail below.

4.1 Closed Form Evaluation Method

The closed-form evaluation method relies on the relationship between location event CFS probability and the observed CFS counts over a geographic area. (Huddleston and Brown 2013) demonstrate that criminal hot-spot (probability) maps can be used to accurately forecast future crime counts within police patrol districts. These criminal hot-spot maps are two dimensional probability density functions that can be estimated using kernel density estimation (Harris 1999), predictive crime models (Smith and Brown 2007; Huddleston and Brown 2009), or by binning historical crime counts by atom (Zhang and Brown 2013). We use the binning approach in this paper. These criminal hot-spot maps provide estimates for the probability of crime occurrence within each atom, with the notation π_i . For district k , the workload score W_k is estimated as the sum of atom event probabilities π_i within the district.

$$W_k = \sum_{i \in k} \pi_i \quad (1)$$

The district workload score W_k represents the proportion of work each district patrol is expected to perform. Since the objective is to provide equal workloads across the districts, district plans are scored using the sample standard deviation of the district workload scores σ_{W_k} . Lower workload standard deviation scores equate to better performance.

Criminal hot-spot maps can also be used to estimate the response time for officers to service calls within their districts. The response time score R is calculated as the sum of the probability weighted distances between each district centroid C_k and each atom location i .

$$R = \sum_{k=1}^K \left[\sum_{i \in k} (\pi_i \|C_k - i\|) \right] \quad (2)$$

In the formula above, the notation $\|C_k - i\|$ denotes the norm (distance) between the district centroid and atom i . Depending on the situation, Euclidean, Manhattan, or travel (road) distance can be used to estimate the travel cost. In this application, we used the Euclidean distance because it offers a simple and fast approximation for the travel distance, especially given that the road structure within Charlottesville is not blocked as in Manhattan. Lower response time scores equate to better district plan performance. This method assumes that there will be very limited cross-boundary service by the patrols within the sectors.

4.2 Discrete-Event Simulation Model Method

The discrete event simulation model is based on the Hypercube Queuing Model (HQM), a well-known descriptive model used to analyze emergency response systems as spatially distributed queuing systems (Larson 1974). In the HQM model, each server (patrol car, fire engine, ambulance, etc.) has two states: idle (0) and busy (1). The state of the whole system is represented as a binary sequence of server statuses. When the number of servers exceeds three, all possible system states form a hypercube.

Historical CFS incidents data and traffic information provide estimates for the arrival rates of the servers into each geographical atom. As long as the aggregate service rate of the system exceeds the total arrival rates of CFS incidents (i.e., supply exceeds the demand), calculating the steady-state probability of the resulting Markov chain provides the probability of being in each possible system state in the hypercube. System performance metrics such as average response time and workload variation are calculated from the hypercube probabilities. While the basic HQM provides a very flexible framework for modeling emergency response systems, the size of the problem grows exponentially with the number of servers. Solving each instance requires solving a linear system with an exponential number of variables (Boyaci and Geroliminis 2011).

Boyaci and Geroliminis (2011) demonstrate that Monte Carlo discrete-event simulations based on HQM converge to the steady-state probabilities estimated by HQM very quickly. Therefore, discrete event simulations provide an alternative method for solving for the HQM steady-state probabilities. The discrete-event model can more easily be extended to simulate complex situations, such as multiple cars responding, different priorities of CFS incidents, different CFS arrival rate at different times of day, as well as various patrol and dispatch rules. We developed the simulation model for the CPD districts in Java 1.6 SE using pseudocode provided by Boyaci and Geroliminis (2011). Sacks (2003) provides the method we used to calculate the expected locations of CFS and patrol cars in the city using the Charlottesville data.

The inputs for the discrete event simulation model are:

- CFS Inter-Arrival Time
- Service Time (Time on Scene)
- CFS Probability for Each Atom
- Geographical Information (District Plan, District Plan Centroids, and Atom Centroids)
- Responding Speed

The simulation model generates CFS using the exponential distribution model defined by the inter-arrival rate parameter. The CFS incidents are spatially distributed within the city according to the geographic probability model generated from historical data as previously discussed (i.e., the CFS probabilities for each atom are the same as those used for the closed form evaluation method). CFS service times are randomly selected from the exponential model defined by the service time parameter. The parameters of the discrete event simulation model are then calibrated such that the queuing parameters (total arrival rate and total service rate) of the system match the historical data set. The discrete-event simulation model tracks the occurrence of four types of events:

- Calls For Service (CFS)
- Patrol Car Arrival at CFS
- Patrol Car Departure from CFS
- Patrol Car Arrival at Base (Idle Position)

When a CFS occurs in the simulation model, the nearest idle patrol car is “dispatched” by changing the server availability status from idle (0) to busy (1). The patrol car (server) status returns to idle once the car returns to its base location within the patrol sector after each event. As for the closed form approach, we use Euclidean distance to approximate travel distance. To simplify the problem, the discrete-event model assumes zero line capacity; if all servers are busy when an incident happens, the incident is “dropped” or considered as being responded to by units outside the modeled system. This mimics the actual police practice; when all cars are busy the CPD requests support from the neighboring Albemarle County Police Department. Dropped incidents do not occur in the lowest intensity setting, rarely occur in the medium intensity setting, and occur less than 4% of the time in the high intensity setting.

The historical Charlottesville Police Department data illustrated in Figure 1 provides estimates for the travel time, service time, and CFS inter-arrival time for the simulation model for each of the three

considered scenarios. The simulation model dynamically tracks the average response time and workload standard deviation measures and stops when these measures converge.

4.3 Agent-Based Simulation Model Method

Agent-based simulation models are increasingly used to model the complex dynamics of resource allocation problems in security applications. Examples include optimizing the location of combat outposts in counter-insurgency (Huddleston, Learmonth, and Fox 2008), examining the use of unmanned surface vehicles for securing Navy ships (Cioppa, Lucas, and Sanchez 2004), and emergency management in disaster response (Wu, Shuman, and Bidanda 2008). Agent-based modeling provides the ability to accurately represent the behaviors of these complex systems by modeling the interactions of the agents of the system.

Zhang and Brown (2013) provide an agent-based model that captures the behaviors of police patrols in a city through the use of model parameters and decision rules. A dynamic view of the agent-based simulation can be seen by clicking the following web address: [agent-based police patrol simulation](#). This model is based on RepastCity from Malleson (2010), which implements agent movement along roads in an urban GIS environment (and provides a flexible framework for changing adapting simulation by changing the associated GIS layers). The simulation model operates by having the agents and environment interact through the use of simple rules. The rules for this agent-based simulation model are:

1. The simulation model generates CFS using the exponential distribution model defined by the inter-arrival rate parameter.
2. CFS incidents are spatially distributed within the city according to the geographic probability model generated from historical data as previously discussed.
3. CFS service times are randomly selected from the exponential model defined by the service time parameter.
4. Police cars randomly patrol the road network within their defined district when not in service.
5. The nearest available patrol car (regardless of district) responds to a CFS at emergency speed.
6. The responding car takes the shortest (road network) path to the location of the CFS.
7. Upon completion of the CFS, if the patrol car is out of its' district, it returns to its district moving at the speed limit and begins patrolling.
8. Upon completion of the CFS, if the patrol car is within its' district, it begins randomly patrolling from its current location.

The input parameters for the agent-based simulation include the CFS inter-arrival time, CFS service time, the emergency speed, and geographic information (the road network, road network speed limits, and the district plan). The agent-based simulation employs the same “dropped call” policy as the discrete event simulation, with similar rates of occurrence. We calibrate the agent-based simulation model by tuning the model until simulating the currently employed districting plan with the simulation model produces the historically observed average response time in all three modeled scenarios. The simulation model dynamically tracks the average response time for all cars and the workload (time in service) proportion for all cars. We run the simulation for a district plan until the average response time and workload proportion for all cars converge to a steady state for the scenario (low-intensity setting, medium intensity setting, or high-intensity setting).

5 DISTRICT PLAN SELECTION

The three district plan evaluation methods above each provide an approach for scoring district plan response time and workload variation performance. The goal is to identify the district plans that provide good performance in both objectives. For multi-objective problems such as this one, there usually does not exist a single solution that simultaneously optimizes both objectives. Instead, there exists a (possibly infinite) set of Pareto-efficient solutions. A solution is Pareto-efficient (also called non-dominated or

Pareto-optimal) for two objectives if one cannot improve performance in one performance measure by selecting a different alternative without sacrificing performance in another. Graphing the performance of the solutions provides a simple way to identify the Pareto-efficient frontier (Gass and Saaty 1955). Figure 3 illustrates the trade-off space and resulting Pareto-efficient solutions identified by the agent-based model in both the low-intensity and high-intensity settings. The agent-based model identifies five non-dominated solutions for the low-intensity scenario and two non-dominated solutions for the high-intensity setting. We used the same approach to identify the non-dominated solution set in each scenario for the closed-form and discrete-event simulations. Figure 3 also identifies these solutions as well as the current district plan used by the CPD.

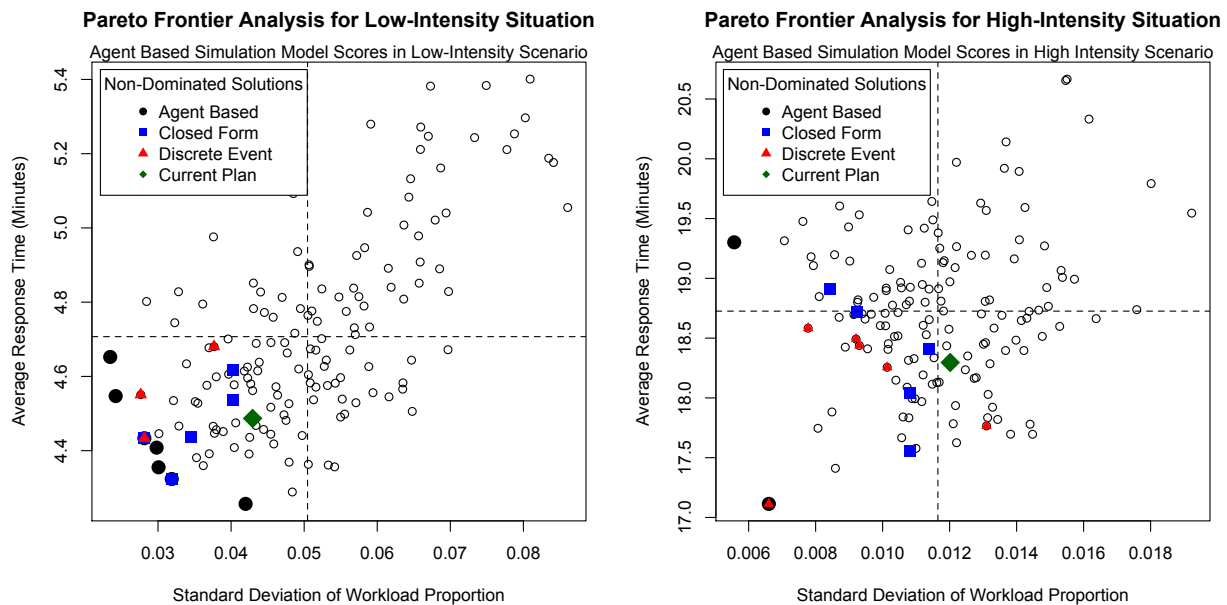


Figure 3: Pareto analysis using the scores from the agent-based simulation in the low-intensity scenario (left panel) and the high-intensity scenario (right panel). Black circles represent non-dominated solutions identified by the agent-based model while colored points represent the non-dominated solutions identified by the other approaches. The dashed lines provide a visual reference for average performance in each evaluation measure as defined by the agent-based model. Note that all three methods provide highly scored non-dominated solutions in the low-intensity scenario but that solutions recommended by the CF and DE methods are rated as relatively average by the agent-based method in the high-intensity scenario. Also note that the average response times on these graphs for the current district plan correlate closely to the observed historical response times for the two scenarios in Figure 1, indicating a well-calibrated simulation model.

6 RESULTS

As Figure 3 illustrates, there is some disagreement between the three different methodologies about which plans are best. Note that in the low-intensity scenario, the agent-based simulation model scores all of the non-dominated solutions by the other two methods relatively highly (they are all clustered in the lower-left hand corner). There is also some agreement on plans that are Pareto-optimal, with some districting plans on the Pareto frontiers of all three methods in the low-intensity scenario. However, in the high-intensity scenario, the Pareto-efficient plans identified by the closed-form and discrete-event methods tend to be rated as relatively average in at least one measure by the agent-based simulation.

Table 1 and Figure 4 provide an explanation for these differences. Table 1 provides the coefficient of determination (R^2) statistic comparing Closed Form (CF), Discrete Event (DE) and Agent-Based (AB) scores for workload variation and response time under low, medium, and high event intensity conditions. Figure 4 provides a pair-wise scatterplot for the most correlated (workload variation in the low-intensity scenario) and least correlated (workload variation in the high-intensity scenario) situations in this table for a visual reference. All methods provide highly correlated workload variation scores in the low intensity scenario but highly uncorrelated scores for workload variation in the high-intensity scenario. The response time scores are less correlated than workload variation scores in the low-intensity scenario but more correlated than workload variation in the high-intensity scenario. For both performance measures, correlation between methods decreases as event intensity increases.

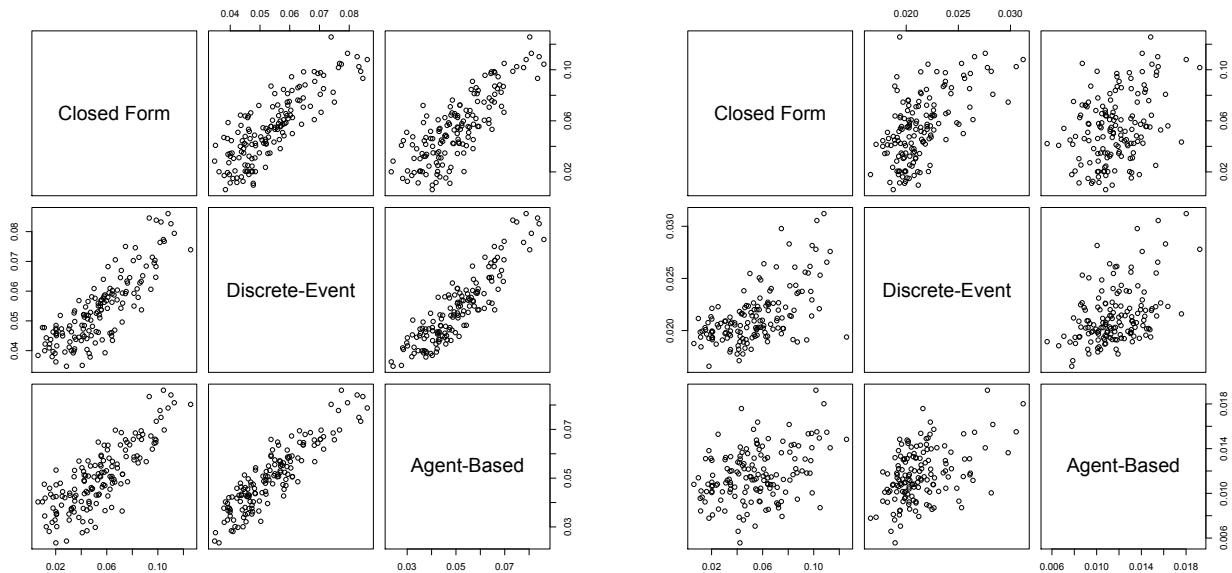


Figure 4: Pairs plots showing pairwise correlation between workload variation scores when event intensity is low (left panel) and when event intensity is high (right panel). All three methods score district plans similarly when event intensity is low but differently when event intensity is high.

Table 1: Table of pairwise R^2 statistics comparing Closed Form (CF), Discrete Event (DE) and Agent-Based (AB) scores for Workload Variation and Response Time under low, medium, and high event intensity conditions.

Event Intensity	Workload Variation			Response Time		
	CF- AB	CF - DE	DE - AB	CF - AB	CF - DE	DE - AB
Low	0.65	0.70	0.81	0.45	0.51	0.74
Medium	0.33	0.44	0.47	0.30	0.44	0.72
High	0.13	0.35	0.19	0.21	0.43	0.52

Table 2 further explains the results observed in Table 1. Table 2 provides the within-method coefficient of determination (R^2) across the three scenarios. As can be seen, the closed form method provides the exact same scores for every scenario, the discrete event scenarios are highly correlated across scenarios, but the agent-based scores change significantly.

Table 2: Table of within-method pairwise R^2 statistics showing how the Workload Variation and Response Time evaluation scores correlate within methods across the three scenarios.

Pairwise Comparison	Workload Variation			Response Time		
	CF	DE	AB	CF	DE	AB
Low-Med	1	0.69	0.45	1	0.92	0.55
Low - High	1	0.60	0.24	1	0.91	0.41
Med - High	1	0.91	0.11	1	0.98	0.46

These results prompted further analysis to understand why the agent-based simulation model scores change so significantly in the high-intensity scenario. We identified two dynamics within the system that cause the agent-based simulation model to significantly alter the scores as CFS intensity increases. The first insight the agent-based model provides concerns the effect of the patrolling behavior of the police cars. When the police cars randomly patrol within their districts, they are often far from the patrol district centroid (as can be seen in the snapshot of the agent-based simulation model in Figure 2). The CPD always dispatches the nearest available police car to the scene of a CFS. Thus, cross-boundary support is quite frequent. In the low-intensity scenario, cross-boundary response averages about 42%. However, this cross-boundary support rises to 70% in the medium intensity scenario and 75% in the high-intensity scenario. These rates roughly correspond to the rates observed in a 1971 New York City study that found that cross-boundary support accounted for more than half of police dispatches (Larson 1971).

The second significant system dynamic is the effect CFS intensity and slow response times due to traffic have on the workload variation during the busy periods of the day. The difference in workload variation (standard deviation of the workload proportion) among districting plans during the high-intensity period around the morning rush hour is very low (note the difference in scales on the horizontal axis in Figure 3). During the high-demand period, all police cars experience a high workload due to the high CFS intensity and slow response speeds due to traffic. Thus, the districting plan has little to do with the workload officers experience during this busy time; for the most part, the police cars are all responding to CFS.

This observation yields an important insight for the CPD. Counter-intuitively, the districting plan becomes most relevant when CFS intensity is low and less important when CFS intensity is high. This is because when CFS intensity is low, the officers spend most of their time patrolling, but when CFS intensity is high, all officers are responding to calls rather than patrolling (on average, 80% of available officer man-hours are employed responding to calls during this period). During the peak rush-hour periods, it may be possible to significantly reduce the average CFS response time by positioning police cars throughout the city near those locations most likely to need CFS during this busy time instead of having officers attempt to both patrol throughout the districts and respond to calls, especially since officers spend relatively little time patrolling the districts they are assigned. In discussions with the CPD, they verified this effect and commented that the system dynamics observed in the agent-based model seemed to correspond closely to that experienced by their officers. In this case, the agent-based simulation model reveals complexities in behavior and applicable insights that the other two evaluation methods do not provide.

7 CONCLUSIONS

Our results indicate that all three evaluation methods produce very similar scores for workload variation when CFS intensity is low enough that the car patrols can meet the demand in their own sectors. However, when the in-district demand exceeds in-district supply, police patrols begin crossing boundaries to meet demand in other police sectors at a very high frequency. This scenario produces a level of complexity that the closed form and discrete event approaches are not well-equipped to handle. Only the agent-based simulation model accurately represents the resulting complexities and significantly changes the workload variation scores to reflect the behavior of the system. The significant insight the agent-based model provides is that, because call volume is so high, officers rarely patrol their sectors in this period, instead spending most (on average about 80%) of their time responding to calls both in and out of sector. The visualization

of the system's complexities the agent-based model provides was also helpful in validating the performance of the simulation with the CPD client.

The scores the three methods provide for response time were less correlated with each other in the low intensity setting than they were for workload variation (ranging between 0.45 and 0.74). However, the correlations between response time scores for the three methods were less sensitive to changes in intensity than the workload variation scores, and the discrete event and agent-based simulations maintained relatively high correlation with each other throughout all three scenarios. The closed form approach did not seem to provide good estimates as it did not have high correlation with either of the other two methods in any of the scenarios. This is probably due to the fact that this method does not account for cross boundary support, and therefore underestimates the effect out of sector CFS have on the average response time.

Future work for this study includes extending the discrete event and agent-based simulation models to dynamically change the modeling parameters for response speed, service time, and inter-arrival time over the 24 hour cycle to correspond with the rates seen in Figure 1. Using this approach will provide an estimate for how well the various districting plans perform over a 24 hour period in actual practice. Planned extensions to the current simulation models include more complex response rules such as call prioritization and multiple car response for certain types of calls. Planned extensions to the closed form approach include performance comparisons using other distance measures (i.e., road-network distance, Manhattan distance, etc.) and development of methods for estimating the effect of cross-boundary support on the average response time performance measure.

REFERENCES

- Altman, M. 1997. "The Computational Complexity of Automated Redistricting: Is Automation the Answer?". *Rutgers Computer Technology Law Journal* 23:81–142.
- Boyaci, B., and N. Geroliminis. 2011. "Extended Hypercube Models for Large Scale Spatial Queuing Systems". Technical report, Ecole Polytechnique Federale de Lausanne.
- Cioppa, T. M., T. W. Lucas, and S. M. Sanchez. 2004. "Military Applications of Agent-Based Simulations". In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 171–180. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Curtin, K. M., K. Hayslett-McCall, and F. Qiu. 2010. "Determining Optimal Police Patrol Areas with Maximal Covering and Backup Covering Location Models". *Networks and Spatial Economics* 10:125–145.
- D'Amico, S. J., S.-J. Wang, R. Batta, and C. M. Rump. 2002. "A Simulated Annealing Approach to Police District Design". *Computers & Operations Research* 29 (6): 667–684.
- Gass, S., and T. Saaty. 1955. "The Computation Algorithm for the Parametric Objective Function". *Naval Research Logistics Quarterly* 2 (1): 39–45.
- Hale, C. D. 1980. *Police Patrol, Operations, and Management*. Upper Saddle River, New Jersey: Prentice-Hall, Inc.
- Hancock, P. G., and N. C. Simpson. 2009. "Fifty Years of Operational Research and Emergency Response". *Journal of the Operational Research Society* 60 (51): 126–139.
- Harris, K. 1999. "Mapping Crime: Principle and Practice". Technical report, National Institute of Justice.
- Huddleston, S. H., and D. E. Brown. 2009. "A Statistical Threat Assessment". *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 39 (6): 1307–1315.
- Huddleston, S. H., and D. E. Brown. 2013. "Using Discrete Event Simulation to Evaluate Time Series Forecasting Methods for Security Applications". In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S. H. Kim, A. Tolk, R. Hill, and M. E. Kuhl. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc. Accepted for publication.
- Huddleston, S. H., G. P. Learmonth, and J. Fox. 2008. "Changing Knives into Spoons". In *Proceedings of the 2008 IEEE Systems and Information Engineering Design Symposium*, 255–260.

- Johnson, D. S. 1985. "The NP-Completeness Column: An Ongoing Guide". *Journal of Algorithms* 6:434–451.
- Larson, R. C. 1971. "Measuring the Response Patterns of New York City Police Patrol Cars". Technical Report R-673-NYC/HUD, RAND Institute, New York City.
- Larson, R. C. 1974. "A Hypercube Queuing Model for Facility Location and Redistricting in Urban Emergency Services". *Computers & Operations Research* 1 (1): 67–95.
- Malleson, N. 2010. *Agent-Based Modelling of Burglary*. Ph.D. thesis, School of Geography, University of Leeds.
- Mitchell, P. S. 1972. "Optimal Selection of Police Patrol Beats". *The Journal of Criminal Law* 63 (4): 577–584.
- Sacks, S. 2003. "Evaluation of Police Patrol Patterns". Technical report, Economics Working Papers, University of Connecticut.
- Smith, M. A., and D. E. Brown. 2007, June. "Application of Discrete Choice Analysis to Attack Point Patterns". *Information Systems and e-Business Management* 5 (3): 255–274.
- Taylor, P. E., and S. J. Huxley. 1989. "A Break from Tradition for the San Francisco Police: Patrol Officer Scheduling Using an Optimization-Based Decision Support System". *Interfaces* 19 (1): 4–24.
- Wu, S., L. Shuman, and B. Bidanda. 2008. "Agent-Based Discrete Event Simulation Modeling for Disaster Response". In *Proceedings of the 2008 Industrial Engineering Research Conference*, 1908–1913.
- Zhang, Y., and D. E. Brown. 2013. "Police Patrol Districting Method and Simulation Evaluation Using Agent-Based Model & GIS". *Security Informatics* 2 (7): 1–13.

AUTHOR BIOGRAPHIES

YUE ZHANG is a Ph.D. Candidate at the University of Virginia, USA. He holds an M.S. from the State University of New York in Systems Science and Industrial Engineering. His email address is yz5yf@virginia.edu.

SAMUEL H. HUDDLESTON is a Ph.D. Candidate at the University of Virginia, USA. He holds an M.S. from the University of Virginia in Systems and Information Engineering. His email address is shh4m@virginia.edu.

DONALD E. BROWN is a Professor in the Department of Systems and Information Engineering, University of Virginia, USA. He serves on the National Research Council Committee on Transportation Security and has served on the National Academy of Sciences panel on High Performance Computing and Crisis Management and on the NRC Committee on Surface Transportation Infrastructure Security. He is a past member of the Joint Directors of Laboratories Group on Data Fusion and a former Fellow at the National Institute of Justice Crime Mapping Research Center. Dr. Brown is a Fellow of the IEEE and a past President of the IEEE Systems, Man, and Cybernetics Society. He holds a Ph.D. from the University of Michigan. His email address is deb@virginia.edu.

GERARD P. LEARMONTH Gerard P. Learmonth Sr. is a Research Associate Professor in the Department of Systems and Information Engineering at the University of Virginia, USA. He holds a secondary appointment as Associate Professor in the Department of Public Health Sciences in the School of Medicine. He holds an M.B.A from New York University, an M.S. in Operations Research from the Naval Post-Graduate School, and a Ph.D. in Statistics and Management Science from the University of Michigan. His email address is jl5c@virginia.edu.