

# Data-Driven Optimization for Atlanta Police Zone Design

Shixiang Zhu, He Wang, Yao Xie\*

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332  
shixiang.zhu@gatech.edu, he.wang@isye.gatech.edu, yao.xie@isye.gatech.edu \*corresponding author

We present a data-driven optimization framework for redesigning police patrol zones in an urban environment. The objectives are to rebalance police workload among geographical areas and to reduce response time to emergency calls. We develop a stochastic model for police emergency response by integrating multiple data sources, including police incidents reports, demographic surveys, and traffic data. Using this stochastic model, we optimize zone redesign plans using mixed-integer linear programming. Our proposed design was implemented by the Atlanta Police Department in March 2019. By analyzing data before and after the zone redesign, we show that the new design has reduced the response time to high-priority 911 calls by 5.8% and the imbalance of police workload among different zones by 43%.

*Key words:* police operations, data analytics, queueing model, optimization

---

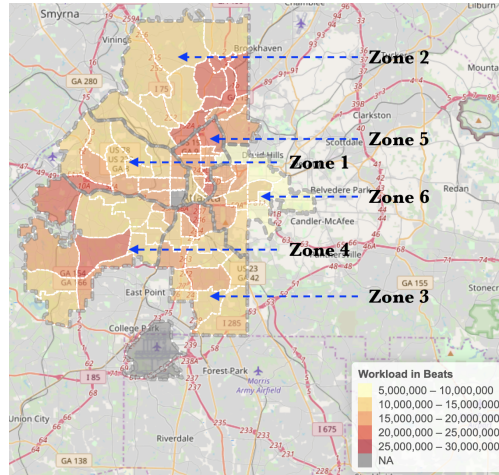
## Introduction

In large urban areas, police departments often organize their patrol forces by dividing the geographical region of a city into multiple patrol areas called *zones* (or *precincts*), and each zone is further divided into smaller areas called *beats* (or *sectors*) (Larson 1972). The design of patrol zones affects both the demand and the capacity for police services in each zone and beat, as well as the travel time of patrol units—together, these factors will determine the police’s response time to emergency calls and crime events. Therefore, the design of patrol zones has a critical impact on the efficiency of police operations.

In this paper, we propose a data-driven framework for designing police zones. The work is developed based on our collaboration with the Atlanta Police Department (APD) for redesigning police zones in Atlanta through a project that lasted from 2017 to 2019. The APD is the primary police force in metro Atlanta, the ninth-largest metropolitan area in the United States and home to 6 million people. The APD divides the geographical area of Atlanta into six zones and 81 beats (see Figure 1). One police patrol unit is usually assigned to each beat to patrol that area and respond to 911 calls. If the response unit is

busy handling another incident, available patrol units in other beats of the same zone will be dispatched to answer calls. Therefore, all patrol units in the same zone can back up for each other, but they typically do not travel across the zone boundaries unless there is a major incident. Under this dispatching policy, each zone can be modeled as a spatially distributed queueing system (Larson 1974a, Larson and Odoni 1981).

**Figure 1** Atlanta Police Zone and Beat Design (Prior to the Redesign in March 2019)



*Notes.* Gray dashed lines represent the boundaries of the six zones. White lines represent the boundaries of the beats. The color in each beat indicates the average police workload level in 2017 (in seconds).

The previous zone and beat configuration in Atlanta (Figure 1) was designed in 2011. Since then, Atlanta has experienced significant population growth and urban development. The U.S. Census Bureau estimated that the population of Atlanta has increased by 15.8% from 2010 to 2017 (United States Census Bureau 2019). The population growth led to an increase in police workload, which was exacerbated by the difficulty faced by the APD in officer recruitment and retention. In 2018, the average response time to high-priority 911 calls (e.g., violent crimes) in Atlanta was 9.5 minutes, which was above the national average (Fritz 2018). Moreover, changing demographics and traffic patterns in Atlanta created uneven police workloads among different regions. The colors in Figure 1 indicate the average workload in each beat, which we compiled from the 911 police incident report data in 2017. A higher workload in a beat often led to longer 911 call response time. For example, Zone 2, which is located in North Atlanta, had a higher than average workload due to its recent commercial development. From the police report data, we found that the average response time for high-priority calls (e.g., carjacks or burglaries) was 13.5

minutes in Zone 2, whereas the average response time citywide for high-priority calls was 9.5 minutes. During the same period, Zone 2, a historically low-crime area, has experienced a rising number of car thefts, burglaries, and armed robberies (Habersham and Deere 2019).


The challenge faced by the Atlanta Police Department motivated us to develop a rigorous quantitative method using operations research for redesigning police zones. The general methods we proposed can also be applied to other urban regions facing similar problems.

## Background of Police Operations

Police departments usually divide the geographical areas of a city into several *zones*, and furthermore each zone is comprised of several *beats*. For example, Atlanta is currently divided into 6 zones and 81 beats (Atlanta Police Department 2020). New York City, which has the largest police department in the U.S., contains 77 precincts (comparable to zones) and 302 sectors (comparable to beats) (New York City Government 2020). Many cities, including Atlanta and New York City, have fixed patrol districts, although in some other cities the districting design may be changed during a 24-hour period if there are significant changes in the pattern of calls throughout the day. We assume the zone districting is fixed over time in this paper.

Typically, one police patrol unit (e.g., a patrol car) is assigned to each beat, where the unit has primary responsibility. When a patrol unit is not busy serving any active calls, it traverses its home beat to perform preventive patrol. If an emergency call is received, the dispatcher will try to send an available unit in the zone to the location of the reported incident. As a result, a patrol unit may be dispatched to any location within its zone, possibly *outside* its home beat. If all units in that zone are busy, the call must wait in a queue until a patrol unit becomes available later. We treat each zone as an independent system since in practice, there are very few dispatches across zone boundaries, due to both long travel time and administrative difficulty (Larson 1972). In line with practice, we also make several assumptions on the dispatching rule summarized in Appendix A.

The typical timeline of a reported incident is as follows (see an illustration in Figure 2): When a 911 call arrives, the dispatcher will try to assign an available patrol unit in the zone to process that incident. If all patrol units are busy, the request needs to wait in a queue until at least one unit in that zone becomes available. We refer to the duration

**Figure 2** The Typical Timeline for Police to Process a 911 Call Incident

between the call arrival time and the dispatch time as the *waiting time*. Once dispatched, the patrol unit travels from its current location to the location of the reported incident. The *response time* is thus the sum of the waiting time and the travel time. After arriving at the scene, the amount of time spent on processing the incident is referred to as the *on-scene time*. The *service time* for the patrol units to process that incident, which we use to calculate the workload, is equal to the sum of the *travel time* and the *on-scene time*.

## Relation to Existing Literature

The study of police zone design and police patrol operations has an extensive history dating back to the 1960s. The books by Larson (1972), Ríos-Mercado (2020) and the survey papers by Chaiken and Larson (1972), Green and Kolesar (2004) provide a comprehensive summary of this field. Below, we review the related literature from three aspects: crime prediction, stochastic modeling of police operations, and optimization for police patrol zones.

*Prediction.* Big data and analytics techniques have gain popularity in policing over the last two decades. Several large police departments in the U.S. have been experimenting with predictive policing that uses historical crime data to predict future crime activities (Perry et al. 2013). For example, Levine et al. (2017) reported the implementation of a predictive policing system by the New York City Police Department. However, a downside of predictive policing is that it does not *prescribe* how predictive crime forecasts should be translated into police’s operational decisions. Saunders et al. (2016) analyzed the predictive policing program of the Chicago Police Department using a quasi-experiment and found that the effect of predictive policing was not statistically significant. The practice of predictive policing has also drawn criticisms from civil rights groups for proliferating racial profiling (Edwards 2016).

*Stochastic Modeling.* Police emergency response systems can be viewed as queueing systems where servers and customers are distributed in space, and servers must travel to



customer locations. Larson (1974a) proposed a hypercube queueing model for urban emergency services. This hypercube queue model is also studied in Chelst and Jarvis (1979), Larson and Odoni (1981). Bammi (1975) considered a beat allocation model where the beat design is fixed, but the patrol units can be moved or shared among different beats.

*Optimization.* Gass (1968) is one of the earliest works that study optimal police beat allocation using integer programming. Mitchell (1972) studied an optimal beat selection problem. Bodily (1978) considered fairness issues of police zone design. Benveniste (1985) proposed a stochastic optimization model that combines zoning and facility location. More recently, Curtin et al. (2005, 2010) proposed maximal covering models for police patrol area design. Cheung et al. (2015), Chow et al. (2015) considered facility location models for police operations. Bucarey et al. (2015) formulated the police districting problem as an enriched  $p$ -median model. Camacho-Collados et al. (2015), Liberatore and Camacho-Collados (2016) introduced a multi-criteria police districting problem that provides a balance between efficiency and workload homogeneity. Piyadasun et al. (2017) considered redesigning police patrol beats by clustering methods. Chen et al. (2019) proposed a street-level design for urban police districting. Implementations of police zone redesign have been reported for Boston, MA (Larson 1974b), Buffalo, NY (Sarac et al. 1999), and Tucson, AZ (Kistler 2009).

Unlike most of the existing literature that focuses on one of the three aspects above, our paper presents a comprehensive framework that *integrates* prediction, stochastic modeling, and optimization for the police zone design problem. We apply this framework to police zone redesign in Atlanta and evaluate the effect of the redesign.

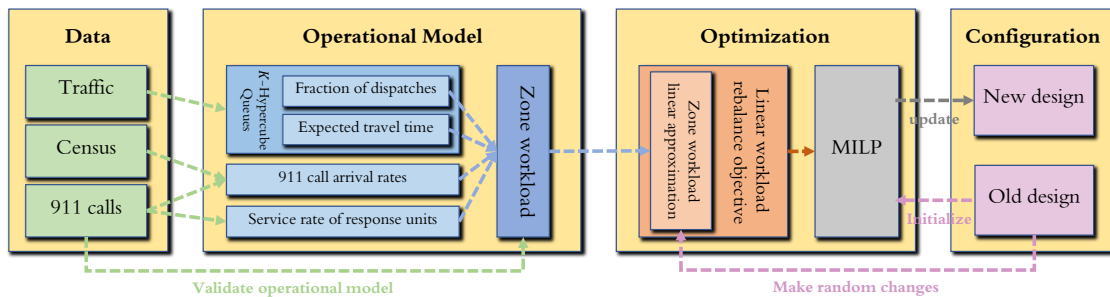
Aside from the police zone design problem, other types of geographical districting problems have also been studied in the operations research literature, including *political districting* (Hess et al. 1965, Garfinkel and Nemhauser 1970) and *school districting* (Ferland and Gu  nette 1990). Several papers (Mills 1967, Morrill 1973, 1976, Vickrey 1961, D’Amico et al. 2002) apply meta-heuristics for geographical districting. These studies consider design criteria such as contiguity (Grofman 1985, Mills 1967, Garfinkel and Nemhauser 1970, Nagel 1972, Mehrotra et al. 1998, Vickrey 1961) and compactness (Garfinkel and Nemhauser 1970, Niemi et al. 1990, Yong 1988). In this paper, we use a network flow based formulation by Shirabe (2009) to impose contiguity constraints. However, the political districting problem and the school districting problem have different objectives than police

districting. Political districting and school districting are often formulated as deterministic optimization problems, but police districting must take into account uncertainty associated with crime occurrence and police service time. Specifically, our model considers a police patrol system that has multiple moving servers (i.e., patrol units). The workload for each server is stochastic, and the servers' status depends on each other. In addition, changes to the zone configuration will impact the operational dynamics of the zone, including the travel time of servers and the number of 911 calls that need to be served.

### Proposed Method: Data Driven Police Zone Design

We propose a data-driven modeling and optimization framework for redesigning police patrol zones. This framework integrates statistical prediction for emergency call arrivals, stochastic modeling of police dispatching process, and optimization for zone redesign in an end-to-end design process. The methodological contribution of this paper is threefold: (a) We create a stochastic queueing system for police operations in Atlanta; the system consists of six interdependent queueing models (each model corresponds to one particular zone), which captures the intra-zone dynamics and the inter-zone effects when the zone design is changed. (b) The parameters of the stochastic system are estimated using real police data; the arrival rates of 911 calls across the city are modeled by a spatio-temporal process. (c) To optimize the zone design using this stochastic model, we overcome the computational challenge by approximating the objective using simulation optimization and finding locally optimal solutions by heuristic search. A summary of this framework is illustrated in Figure 3. Next, we describe each component of the framework.

**Figure 3** An Illustration for the Data-Driven Optimization Framework of Police Zone redesign




## Module I. Stochastic Model for Police Patrol and Emergency Response

The basis of our design process is a stochastic model for the police's response to 911 calls and processing time. For a *fixed* zone configuration, the dynamics of patrol units in each zone are independent, since we assume patrol units do not respond to requests from a different zone than they are assigned. Within each zone, the movement and workload of patrol units are quite complex, as they are determined jointly by the zone design, the dispatch rule, and the beat-to-beat travel time.

We model emergency response operations based on the *hypercube queue model* (Larson 1974a); see Appendix B for details. The system state depends on the status of all the patrol units in this zone and the number of calls in the queue to be processed. Consider the state as a vector of binary numbers, where each element of the vector represents the status of the corresponding unit: the unit is busy processing a call if the corresponding digit of the state equals 1 and otherwise 0. Based on whether there is any available unit, the state space can be divided into two parts: (a) *unsaturated states*: the states where the queue for unprocessed calls is empty. These states can be represented by a hypercube, where each vertex corresponds to a state. (b) *saturated states*: the states where the queue for unprocessed calls is non-empty. By definition, all patrol units are busy under saturated states. An example with three patrol units is shown in Figure 4.

**Figure 4** The State Space of a Hypercube Queue Model Representing a Zone with Three Beats



*Notes.* Nodes  $B_0, \dots, B_7$  represent unsaturated states (at least one patrol unit is available) and nodes  $S_1, S_2 \dots$  represent saturated states (all patrol units are busy).

We define the system dynamics using a transition rate matrix based on call arrival rates, travel time, and service rates. There are two classes of transitions on the hypercube: *upward*

*transitions* that change a unit's status from idle to busy; *downward transitions* that do the reverse. The downward transition rate from one state to its adjacent states is always the service rate. The upward transition rates, however, will depend on the dispatch rule when a call is received and the travel time between beats. We assume that when a call is ready to be processed, the nearest available unit with the minimum mean travel time to the location of the call is dispatched.

The steady state probabilities of this system define the likelihood that each state is occupied in the long run, which can be determined by solving the balance equations. We note that although the number of saturated states is infinite, it can be expressed in closed-form as in an M/M/c queue. The main challenge for computing the steady state distribution is that the number of *unsaturated* states on the hypercube grows exponentially with the number of beats. We developed an efficient computation method to approximate the steady state distribution by exploiting the fact that the transition rate matrix is sparse.

## Module II. Model Estimation and Prediction


We estimate the parameters of the hypercube queue model using data provided by the Atlanta Police Department (APD) and public data sources. We first describe the datasets that we used for estimation. Then we describe our call arrival rate and workload prediction method.

**Data Sets.** We have utilized two key data sets for police workload prediction:

1. *911 Call Data*: the APD has electronic records of police reports, where each record consists of detailed information about a 911 call, such as the call time, dispatch time, arrive time, clear time, and location. These reports also provided travel information of each dispatch, including a departure time and location, and an arrival time and location. Using this data set, we estimate the average travel time between beats, the 911 call arrival rates for each beat, and the service rate of patrol units.


2. *American Community Survey*: the U.S. Census Bureau provides an annual American Community Survey (ACS), which includes comprehensive information about the population, demographic, and economic status of different areas in Atlanta. We use population, number of housing unit, school enrollment, median household income, median number of rooms, median age, median house price, and average age of houses from the ACS as explanatory variables in our predictive model.

**Figure 5** Distribution of the On-Scene Time of All Units



*Notes.* On-scene time is approximated by exponential distributions. The red vertical line is the mean on-scene time for all response units.

**Figure 6** Beat-to-Beat Travel Time Matrix Estimated from Real Data



*Notes.* Each row and each column corresponds to a beat, and each small square represents average beat-to-beat travel time.

**Call Arrival Rate Prediction.** We cannot use standard stochastic processes to model 911 call arrivals, because the arrival rates have a significant seasonality pattern and yearly trend; the rate also have a strong spatial correlation among adjacent geographical areas. Therefore, we propose a spatio-temporal model to predict call arrival rates using spatially lagged endogenous regressors (Rosen 1974). We assume that the 911 call arrival rate of a beat depends on (a) the arrival rates of adjacent beats, (b) the arrival rate of the same beat in the previous year, and (c) the demographic factors in the past five years. We also assume that the noise term has a spatial structure, where the covariance between two beats is determined by an exponential kernel function. The parameters of the model are fitted by maximum likelihood estimation. See Appendix C for details.

**On-scene Time Estimation.** As we assume the on-scene time of a response unit is independent of its service region, we analyze the distribution of average on-scene time for all units in the same zone. We have observed from data that the on-scene time of response units fits to an exponential distribution shown in Figure 5.

**Travel Time Estimation.** We estimated the travel time between beats from the 911 call dispatch data. The estimation result is shown in Figure 6.

**Performance Metrics.** Using the steady state distribution of the hypercube queue model with estimated model parameters, we can derive several key performance measures (See Appendix B for details), which will guide our zone design decisions in the next step:


- *Expected Travel Time.* The travel time for a dispatch is affected by both the beat-to-beat travel time and the fraction of dispatches. The fraction of dispatches critically depends on the utilization factor of the hypercube queue system. If the utilization is low, most calls can be processed by units close to their home beats, resulting in a shorter travel time. However, if the system is congested, calls must wait in the queue, which will be processed by the first unit available. This will lead to a longer travel time.
- *Response Time.* The response time for a call includes both the travel time for the patrol unit to arrive at the scene and the possible waiting time if the call incurs a positive queue delay.
- *Patrol Unit Workload.* The workload of one patrol unit is the sum of the steady state probabilities for all the states where this unit is busy. The total workload for all patrol units in the zone is simply the aggregation of individual unit workloads.

**Model Validation.** To evaluate the predictive accuracy of the estimated police zone workload, we compare the real zone workload reported by historical data with the out-of-sample zone workload simulation from our operations model using leave-one-out cross-validation. Specifically, the operations model takes the 911 calls arrival rates, the police service rate, and beat-to-beat travel time in a particular year as input variables and generate the simulated police workload for each zone using the existing zone design. In practice, we estimate these input quantities using the data from the past five years. Figure 7 shows comparisons for the average workload between zones and averaged workload between years. To carry out the evaluation for each zone/year in Figure 7, the model estimation is performed using all data except for one zone/year, and a prediction is made for that zone/year. The result shows that the simulation is consistent with the real data, and the percentage difference is less than 6.7%.

The operations model allows us to quantify the impact of an arbitrary zone design on the zone workload analytically. However, it cannot be used to search for an improved zone design efficiently due to the randomness and the non-linear structures of zone workloads with respect to decisions. In the next section, we will develop a way to approximate the zone workload, which can be used in the optimization.



**Figure 7 Simulated Workload from the Hypercube Queue Model and Real Workload Data for Each Zone (left) and Each Year (right)**



Notes. Data is only available for 2017 Q1.

### Module III. Zone Design Optimization

We develop an optimization model for zone redesign. We assume that the shape of the beats are fixed, so the zone design problem can be formulated as a graph partition problem, where we allocate the beats to a pre-determined number of zones.

Different objectives have been studied in the literature for police zone design (see Appendix D). After several discussions with the Atlanta Police Department, we learned that their primary concern is the workload imbalance among different geographical areas, as they have observed a trend of police overwork in some zones. In particular, longer shifts and police officer fatigue may lead to a host of public safety issues and affect the police department's budget (in the form of bonuses and overtime). Therefore, in this paper, we follow the zone workload definition introduced by Mitchell (1972) and choose the objective function as the workload variance among different zones, which quantitatively measures the police workload imbalance among zones. This metric was also used internally by APD to measure previous zone redistricting plans (Egbert 2016). The optimization model aims to minimize the workload imbalance across zones subject to some shape constraints for each zone, such as contiguity and compactness. The full optimization model is included in Appendix E.

**Zone-level Workload Approximation.** A key challenge for solving the optimization problem is characterizing the complex dependence between the zone design and the zone-level workload. In theory, the expected workload in each zone can be computed using the hypercube queue model given any zone design. However, it is impractical to carry out such computation to search for an improved zone design since the decision space is extremely

large. For example, with the six zones and 81 beats in Atlanta, this will generate more than  $2.4 \times 10^{59}$  possible zone designs, where each zone design corresponds to a separate queueing model.

To tackle this challenge, we consider a local linear approximation of the workload function and then use the approximated value in the objective function of the optimization model, which allows the objective function of the optimization problem to be expressed explicitly without repeatedly computing the hypercube queue model. This linear approximation represents the first-order Taylor expansion of the workload change when the current zone design is modified slightly by adding a few beats or removing a few beats. In Appendix E, we discuss the detailed process of generating this linear approximation and validating the effectiveness of our approximation method using randomly selected out-of-sample designs.

**Contiguity and Compactness Constraints.** In addition to balancing the police workload, the shapes of zones should be contiguous and compact. Since patrol units almost always travel within a zone, a narrow or snakelike shape increases the travel time and reduces the efficiency of patrol operations and emergency response. Therefore, we formulate the *contiguity* and *compactness* criteria using binary variables and a set of linear constraints, which requires that beats in the same zone to be geographically connected and close to each other; see Appendix E for details.

**MILP Formulation and Local Search.** With the linear approximation of the zone workload, the objective function measuring the workload variance becomes a *quadratic* function of the binary decision variables. Using the McCormick envelopes technique in integer programming (McCormick 1976), this quadratic function can be expressed equivalently as a linear form. The reformulated mixed-integer linear programming (MILP) model contains more than 240,000 variables and has a huge solution space with more than  $2.4 \times 10^{59}$  possible solutions.

Based on the feedback from APD, we also consider a practical constraint that beat changes should be as few as possible. A drastic zone design change is undesirable for two reasons: First, a large-scale operational change will result in high implementation costs; second, a radical design change will usually face significant uncertainties and unpredictable risks in the future operation. For this reason, we only consider plans with no more than six beat changes. Moreover, two specific design requirements were suggested by APD and

included as constraints in our model: (a) beat 305, 111 should be staying in zone 3 and 1, respectively; (b) zone 3 will not take any existing beat from zone 4.

In order to find a near optimal solution to the MILP efficiently and avoid limit the number of beat changes from the existing design, we use a local search method based on simulated annealing. The optimization process takes about 30 to 50 minutes on a standard laptop with a quad-core 4.7 GHz processor.

## Implementation Results

During a one-year collaboration between Georgia Tech and APD, we used our data-driven zone design framework as a decision support tool and implemented a zone reconfiguration plan. In the following, we first discuss the police workload imbalance analysis using real and predicted results; then we present our zone redistricting plan and the major changes we made to the original zone configuration; we also discuss the implementation's impact; lastly, we investigate the effectiveness of the proposed plan by analyzing the post-implementation data.

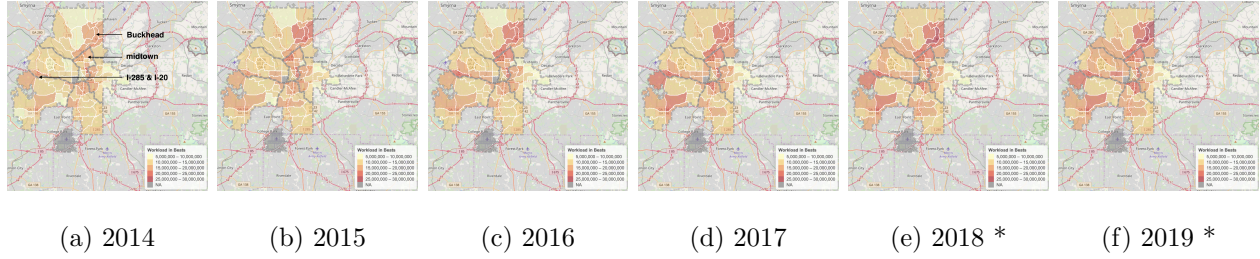
### Workload Imbalance Analysis

Under the previous zone design in Atlanta, the workload imbalance across zones had worsened in the past few years. This imbalance is partly due to the growth and redistribution of population, as well as changing (mostly worsening) traffic conditions. Here, we analyze how zone configuration affects the workload imbalance over the years, and predict the future trend if there is no change in the configuration. To quantify the imbalance of workload, we use the variance of workload across zones as the metric.

We compiled the actual workload from the year 2014 to 2017 (the most recent year at the time of our redesign process) from the police report data, and computed the predicted workload for the year 2018 and 2019 using the stochastic queue model we developed earlier. The workload in each beat is displayed on the city map in Figure 8. One can observe a clear trend that both the mean and the variance of workload were growing up to 2017, and would continue to grow according to our model prediction. Moreover, the workload was exceptionally high in several concentrated areas, such as Buckhead in Zone 2, Midtown in Zone 5, and the intersection of highway I-285 and highway I-20 in Zone 4.


In Figure 9, we plot the total workload of the entire city (red curve) and the workload imbalance across the six zones (blue curve). As before, the results for 2018 and 2019 are

**Figure 8 The Beat Level Workload (in Seconds) in Atlanta from 2014 to 2019 (2018 and 2019 Levels are Predicted by the Hypercube Queue Model)**



*Notes.* Dashed grey lines represent the boundaries of the six zones. Darker color means a higher workload.

**Figure 9 The Total Workload of All Zones (Red) and Variance of Workload Across Zones (Blue) from 2013 to 2019**



*Notes.* The values for 2018 and 2019 are predicted levels assuming the configuration is not changed, and the stars for the year 2018 and 2019 are predicted levels under our proposed redesigned configuration.

based on predictions assuming the same zone design. The figure shows that both the total workload and the workload variance would have increased without zone reconfiguration.

### Zone Reconfiguration

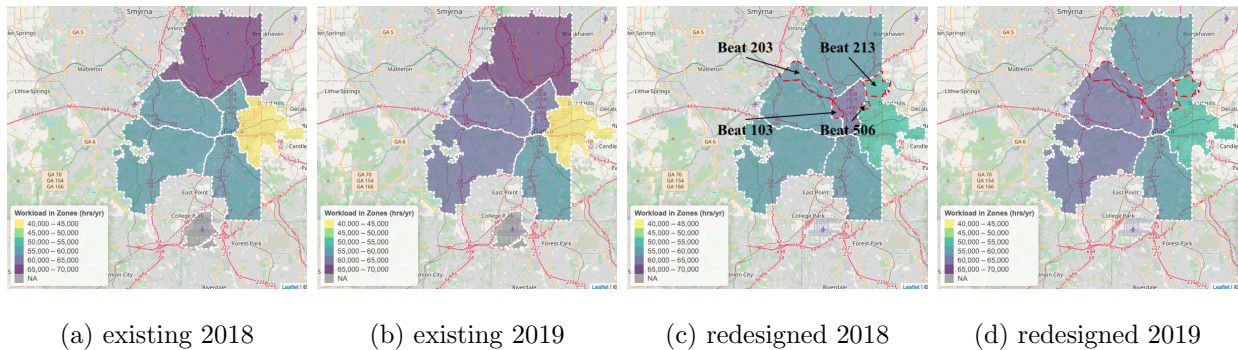
Using the queueing and optimization models developed earlier in this paper, we proposed a new police zone plan for Atlanta. Our analysis report in 2018 contained beat-wise workload prediction for the next two years (2018 and 2019) and proposed three candidate designs with similar beat shifts that all attain improved workload balance. In Table 1, we list the predicted annual workload in each zone, total workload, and workload variance. After our proposed plan was reported to the police, several follow-up meetings were held to elicit feedback from the APD and deliberate various design trade-offs. The APD Deputy Chief and other senior police officers also participated in these discussions and provided

feedback. After several rounds of discussions, the police chose one of our proposed designs as the new police zone configuration. The new design and the previous design are shown in Figure 10. This particular design was selected for three major reasons: (a) this plan would only change four beats, so it has fewer beat shifts in comparison with other candidate plans; (b) Beat 203 and 213 are two areas with relatively low police workload in zone 2, so moving them out of zone 2 may help focus their police resources on curbing the crime surge in the Buckhead area (the center of zone 2); (c) Zone 3 and Zone 4 have maintained a good operational balance for the recent years; keeping these two zones intact can also help reduce the implementation cost of the plan. For both predictions in 2018 and 2019, the redesigned zone plan has a more balanced workload distribution over the six zones (the colors in Figures 10 (c) and 10 (d) are more uniform). The new design is projected to reduce the workload variance significantly by 62.79% and 55.97% in 2018 and 2019, respectively.

**Table 1** Comparison of the Predicted Workload Under the Previous Zone Design and the New Zone Design

	Workload ( $\times 10^4$ hrs/yr)						Total	Variance ( $\times 10^7$ )	Variance increase ratio compare with 2016 (%)
	Zone 1	Zone 2	Zone 3	Zone 4	Zone 5	Zone 6			
Real 2016	5.3744	6.0082	5.3891	5.6164	5.5479	4.1413	32.0776	3.3441	N/A
Real 2017	5.6648	6.2981	5.5154	5.7919	5.7319	4.2454	33.2479	3.9454	+17.98
Predicted 2018	5.9558	6.5833	5.6366	5.9606	5.9143	4.3488	34.3996	4.6375	+38.67
Predicted 2019	6.2487	6.8650	5.7642	6.1292	6.0976	4.4518	35.5568	5.4267	+62.27
Predicted 2018 with redesign	5.9018	5.6961	5.6366	5.9606	6.1595	5.0448	34.3996	1.2442	-62.79
Predicted 2019 with redesign	6.1833	5.9294	5.7642	6.1292	6.3686	5.1819	35.5568	1.4721	-55.97

**Figure 10** Comparison of Zone Workload Distribution Between the Existing Plan and the Redesigned Plan for 2018 and 2019



*Notes.* The color depth of each zone represents the level of its annual workload. Beat shifts in the redesigned plans have been highlighted by red dash lines. As we can see, for the predicted year 2018 and 2019, the redesigned plan achieves a more balanced workload across zones comparing to the existing plan.

## Implementation and Impact

In January 2019, we submitted the final report for zone redesign to the Atlanta Police Department, which the Police Chief and her team reviewed. Our report analyzed the annual growth trend of police workload and proposed a detailed redistricting plan. Our redistricting plan involves changing four zones: Zone 6 in East Atlanta will increase by four square miles. Zone 1 in Northwest Atlanta will increase by two square miles. Zone 2, which covers Northeast Atlanta and Buckhead, decreases by seven square miles. Zone 5 also has some minor changes. The detailed plan is shown in Figure 10. Overall, the redistricting would rebalance the police workload between the four zones and reduced the average response time in Zone 2 (Habersham 2019).

In February 2019, Atlanta’s City Council voted to approve our proposed redistricting plan. The Atlanta Police Department officially implemented the plan on March 17, 2019. In a statement about this redesign (Atlanta Police Department 2019), the APD Deputy Chief Jeff Glazier stated: “It is important that we examine our officer workload periodically, and with the help of Georgia Tech, we were able to do so in a data-driven manner. We are confident these changes will balance the workload in all zones.” The new zone design was also covered by several media outlets, including the Atlanta Journal-Constitution, the largest daily newspaper in the metro Atlanta area. On social media, many residents praised the change and thanked the APD and the Georgia Tech team for the new plan.

## Post-Implementation Analysis

In August 2020, we received a new set of 911 call data provided by the Atlanta Police Department thanks to Major John Quigley’s help. This dataset contains records from the pre-implementation period (March 2017 ~ March 2019) to the post-implementation period (March 2019 ~ March 2020). Note that we excluded the data after March 2020 because police operations had been drastically affected by the COVID-19 pandemic and the citywide lockdown. Also, due to the inconsistency between the new data format and the previous data format, we normalized the new data set by bringing the workload of each zone in the overlapping period (i.e., Year 2017) to the same level. Using this dataset, we evaluate the zone-level workload distribution and variance before and after the implementation. Apart from the workload analysis, we also focus on three other metrics as we described in the background section, namely, waiting time, travel time, and response time. (We do not study on-scene time since it is exogenous to the zone design.) In particular, response time

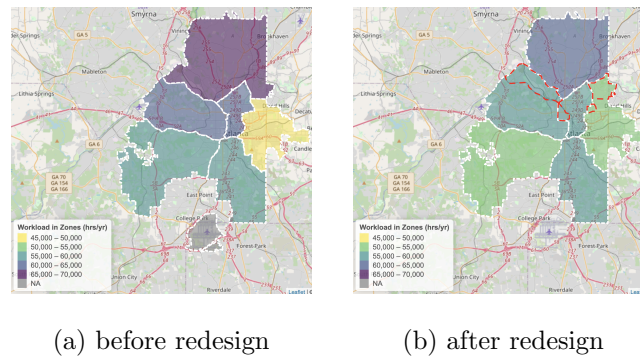


is important for the perception of police operations efficiency by the general public and directly reflects the effectiveness of the districting plan.

**Table 2 Comparison of the Real Workload Before and After the Redesign**

Time period	Workload ( $\times 10^4$ hrs/yr)							Variance ( $\times 10^7$ )	Variance increase year-on-year ratio (%)
	Zone 1	Zone 2	Zone 3	Zone 4	Zone 5	Zone 6	Total		
Mar 2018 ~ Mar 2019 (before)	6.0568	6.8737	5.5450	5.7040	6.1607	4.7360	35.0762	4.2375	+7.40
Mar 2019 ~ Mar 2020 (after)	5.5142	6.3753	5.7157	5.0090	5.9462	5.0129	33.5733	2.3925	-43.54

**Figure 11 Zone-Level Average Workloads Before (Left) and After (Right) the Redesign**



*Notes.* The color depth of each zone represents the level of its annual workload. The beats changed in the redesigned plans (right) are highlighted by red dash lines.

In Table 2 and Figure 11, we show the average workload of each zone and their variances before and after the redesign in March 2019. It is evident that our redesign successfully mitigated the workload increase rate and reduced the deterioration of workload imbalance: the workload variance of 2019 decreased by 43.54% comparing to 2018. For Zone 2, a region with rapidly increasing crime rates and a focus of this redesign, the analysis shows that our new plan has reduced its police workload by 7.25% compared to 2018.


We also report the average response time, waiting time, and travel time for all 911 calls and high-priority calls (e.g., serious crimes) in Tables 4 and 3. All three time metrics have worsened from March 2018 to March 2019, the year before the zone redesign, as a result of economic development and population growth. After the new zone design was implemented, all three metrics have seen decreased growth rates. For high-priority calls, in particular, these time metrics have decreased significantly. In Figure 12, we also show the distributions of the three metrics for high-priority calls; the redesign has resulted in the probability density shifting left toward zero.

**Table 3 Average Response, Waiting, and Travel Time Per Call for High-Priority Calls in Six Zones**

Category	Time period	Time per call (minutes)							Time increase year-on-year ratio (%)
		Zone 1	Zone 2	Zone 3	Zone 4	Zone 5	Zone 6	Citywide	
Response	Mar 2018 ~ Mar 2019 (before)	11.94	11.61	10.91	12.99	13.57	9.81	11.90	+10.90
	Mar 2019 ~ Mar 2020 (after)	13.92	10.59	12.06	9.36	10.34	9.80	11.21	-5.80
Waiting	Mar 2018 ~ Mar 2019 (before)	3.49	3.16	4.69	3.22	3.39	2.46	3.55	+36.02
	Mar 2019 ~ Mar 2020 (after)	4.91	2.38	3.41	3.31	2.95	3.29	3.47	-2.25
Travel	Mar 2018 ~ Mar 2019 (before)	8.48	9.55	17.73	9.45	10.05	7.08	10.72	-0.74
	Mar 2019 ~ Mar 2020 (after)	9.59	8.64	8.28	7.73	6.59	7.36	8.07	-24.72

**Table 4 Average Response, Waiting, and Travel Time Per Call for All Cases in Six Zones**


Category	Time period	Time per call (minutes)							Time increase year-on-year ratio (%)
		Zone 1	Zone 2	Zone 3	Zone 4	Zone 5	Zone 6	Citywide	
Response	Mar 2018 ~ Mar 2019 (before)	30.35	32.21	27.81	30.70	30.42	26.07	29.69	+11.70
	Mar 2019 ~ Mar 2020 (after)	28.72	24.77	24.70	28.59	36.20	33.25	31.26	+5.29
Waiting	Mar 2018 ~ Mar 2019 (before)	18.41	17.75	16.61	17.87	17.97	14.21	17.24	+12.83
	Mar 2019 ~ Mar 2020 (after)	18.00	13.96	14.63	16.56	21.74	19.63	19.08	+10.67
Travel	Mar 2018 ~ Mar 2019 (before)	8.20	7.88	7.49	8.60	6.80	6.27	7.50	+28.21
	Mar 2019 ~ Mar 2020 (after)	7.56	6.02	7.28	7.51	5.98	6.27	7.43	-0.93

**Figure 12 Distribution of Response, Waiting, and Travel Time Per Call for High-Priority Calls Before and After the Redesign, Calculated by Kernel Density Estimation**

Notes. Vertical Dotted Lines Indicate the Mean of Each Distribution.

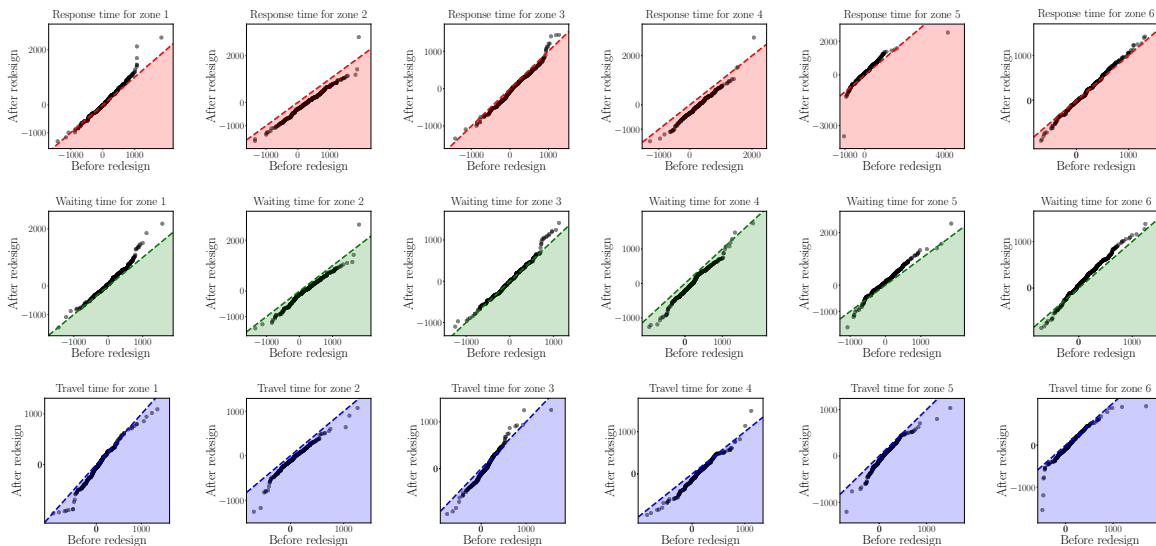
To control for the time series effect, we also perform a *difference in difference* (DID) analysis for the yearly variance change of the three time metrics. We calculate the variance of these metrics on each day of the year and then compute the variance difference between two years. We then compare the variance difference before and after the redesign. In Figure 13, we report this comparison for the entire Atlanta. For travel time, the result shows that the variance difference after the redesign is significantly lower than the difference before the redesign, as most points are below the 45-degree line. For response and waiting time, the figure also shows a slight decrease after the redesign. In Figure 14, we show the

**Figure 13** The difference in difference (DID) analysis of variance for response time (red), waiting time (green), and travel time (blue)



*Notes.* Each point represents one day of the year. The x-axis corresponds to the variance difference for the two years prior to the redesign, and the y-axis corresponds to the variance difference for the two years before and after the redesign. The 45-degree dashed line passes through the origin. Points below the dashed line indicate in a reduction in variance increase.

**Figure 14** Difference in Difference (DID) of Variance for Response Time (1st row), Waiting Time (2nd Row), and Travel Time (3rd Row) in Each Zone



same comparison for each zone separately. For Zones 2, 3, and 4, which have experienced upticks in crime rates and police workloads (see Figure 11), the variance differences for all three metrics have decreased.

## Final Remarks

In this paper, we presented a data-driven framework for police zone redesign. The framework includes three main components: a spatial queueing model for police patrol and emergency response operations, a statistical model to predict emergency call arrivals and police

service rates, and a mixed-integer linear programming model for zone districting optimization. We applied this framework to redesigning police zones in Atlanta. The Atlanta Police Department officially adopted our proposed redesign plan in March 2019. By analyzing data before and after the zone redesign, we show that the new design has reduced the response time to high-priority 911 calls by 5.8% and the imbalance of police workload among different zones by 43%.

We note several potential limitations in our method that can be improved in future research: (a) our model assumes that the occurrence of crime events and 911 calls are exogenous to the police zone design; however, we observe in the post-implementation data that the call rates changed after the zone redesign, which imply that these factors may be endogenous and affected by the zone redesign; (b) our method assumes that beat designs are fixed, and hence the zone design problem can be formulated as an assignment problem. For a complete redesign, beat designs and zone designs can be optimized jointly.

## Acknowledgments

We are grateful to Chief Erika Shields, Major John Quigley, Captain Jacquelyn Gwinn-Villaroel, and other police officers of the Atlanta Police Department for their valuable inputs and assists throughout this project. We would also like to thank Mr. Song Wei for helping in the data preparation. This research is supported in part by the Atlanta Police Foundation and the National Science Foundation (CMMI-2015787).

## References

- Atlanta Police Department (2019) Atlanta police make changes to four zone boundaries as part of beat redesign. <http://www.atlantapd.org/Home/Components/News/News/190/> (accessed May 1, 2020).
- Atlanta Police Department (2020) Atlanta police department zones. <https://www.atlantapd.org/community/apd-zones> (accessed May 1, 2020).
- Bammi D (1975) Allocation of police beats to patrol units to minimize response time to calls for service. *Computers & Operations Research* 2(1):1–12.
- Benveniste R (1985) Solving the combined zoning and location problem for several emergency units. *Journal of the Operational Research Society* 36(5):433–450.
- Bodily SE (1978) Police sector design incorporating preferences of interest groups for equality and efficiency. *Management Science* 24(12):1301–1313.
- Bucarey V, Ordóñez F, Bassaletti E (2015) Shape and balance in police districting. *Applications of Location Analysis*, 329–347 (Springer, Cham, CH).
- Camacho-Collados M, Liberatore F, Angulo JM (2015) A multi-criteria police districting problem for the efficient and effective design of patrol sector. *European journal of operational research* 246(2):674–684.

- Chaiken JM, Larson RC (1972) Methods for allocating urban emergency units: a survey. *Management Science* 19(4-part-2):110–130.
- Chelst K, Jarvis JP (1979) Estimating the probability distribution of travel times for urban emergency service systems. *Operations Research* 27(1):199–204.
- Chen H, Cheng T, Ye X (2019) Designing efficient and balanced police patrol districts on an urban street network. *International Journal of Geographical Information Science* 33(2):269–290.
- Cheung CY, Yoon H, Chow AH (2015) Optimization of police facility deployment with a case study in greater london area. *Journal of Facilities Management* 13(3):229–243.
- Chow AH, Cheung C, Yoon H (2015) Optimization of police facility locationing. *Transportation Research Record* 2528(1):60–68.
- Curtin KM, Hayslett-McCall K, Qiu F (2010) Determining optimal police patrol areas with maximal covering and backup covering location models. *Networks and Spatial Economics* 10(1):125–145.
- Curtin KM, Qui F, Hayslett-McCall K, Bray TM (2005) Integrating gis and maximal covering models to determine optimal police patrol areas. *Geographic information systems and crime analysis*, 214–235 (IGI Global).
- D’Amico SJ, Wang SJ, Batta R, Rump CM (2002) A simulated annealing approach to police district design. *Computers & Operations Research* 29(6):667–684.
- Edwards E (2016) Predictive policing software is more accurate at predicting policing than predicting crime. <https://www.aclu.org/blog/criminal-law-reform/reforming-police/predictive-policing-software-more-accurate-predicting> (accessed May 1, 2020).
- Egbert J (2016) Beat redistricting proposal 2016. Atlanta Police Department (internal report).
- Ferland JA, Guénette G (1990) Decision support system for the school districting problem. *Operations Research* 38(1):15–21.
- Fritz J (2018) What is the average police response time in the U.S.? <https://www.safesmartliving.com/home-security/average-police-response-time/> (accessed April 30, 2019).
- Gaetan C, Guyon X (2010) *Spatial Statistics and Modeling* (Springer, New York).
- Garfinkel RS, Nemhauser GL (1970) Optimal political districting by implicit enumeration techniques. *Management Science* 16(8):B-495–B-508.
- Gass SI (1968) On the division of police districts into patrol beats. *Proceedings of the 1968 23rd ACM National Conference*, 459–473 (ACM).
- Green LV, Kolesar PJ (2004) Anniversary article: Improving emergency responsiveness with management science. *Management Science* 50(8):1001–1014.
- Grofman B (1985) Criteria for redistricting: A social science perspective. *UCLA Law Review* 33(1):77–184.

- Habersham R (2019) Atlanta police hope changes to patrol zones shorten response times. *The Atlanta Journal-Constitution* <https://www.ajc.com/news/local/atlanta-police-change-zoning-boundaries-help-with-response-times/R4hjmqmYEmaveOgkW6NfRU0/>.
- Habersham R, Deere S (2019) Buckhead residents confront mayor, police chief about crime. *The Atlanta Journal-Constitution* <https://www.ajc.com/news/local/buckhead-residents-confront-mayor-police-chief-about-crime/5VYluSvFkIAMqUloAIIRSK/>.
- Hess SW, Weaver J, Siegfeldt H, Whelan J, Zitlau P (1965) Nonpartisan political redistricting by computer. *Operations Research* 13(6):998–1006.
- Kistler A (2009) Tucson police officers redraw division boundaries to balance their workload. *Geography & Public Safety* 1(4):3–5.
- Larson RC (1972) *Urban police patrol analysis* (MIT Press).
- Larson RC (1974a) A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers & Operations Research* 1(1):67–95.
- Larson RC (1974b) Illustrative police sector redesign in district 4 in boston. *Urban Analysis* 2(1):51–91.
- Larson RC, Odoni AR (1981) *Urban Operations Research* (Prentice Hall, NJ).
- Lee ML, Goldsman D, Kim SH, Tsui KL (2014) Spatiotemporal biosurveillance with spatial clusters: control limit approximation and impact of spatial correlation. *IIIE Transactions* 46(8):813–827.
- Levine E, Tisch J, Tasso A, Joy M (2017) The New York City police department’s domain awareness system. *Interfaces* 47(1):70–84.
- Liberatore F, Camacho-Collados M (2016) A comparison of local search methods for the multicriteria police districting problem on graph. *Mathematical Problems in Engineering* 2016(3):1–13.
- McCormick GP (1976) Computability of global solutions to factorable nonconvex programs: Part iconvex underestimating problems. *Mathematical Programming* 10(1):147–175.
- Mehrotra A, Johnson EL, Nemhauser GL (1998) An optimization based heuristic for political districting. *Management Science* 44(8):1100–1114.
- Mills G (1967) The determination of local government electoral boundaries. *Operations Research Quarterly* 18(3):243–255.
- Mitchell PS (1972) Optimal selection of police patrol beats. *Journal of Criminal Law, Criminology & Police Science* 63(4):577–584.
- Morrill RL (1973) Ideal and reality in reapportionment. *Annals of the Association of American Geographers* 63(4):463–477.
- Morrill RL (1976) Redistricting revisited. *Annals of the Association of American Geographers* 66(4):548–553.



- Nagel SS (1972) Computers and the law and politics of redistricting. *Polity* 5(1):77–93.
- New York City Government (2020) NYPD precincts and sectors. <https://data.cityofnewyork.us/Public-Safety/NYPD-Sectors/eizi-ujye> (accessed May 1, 2020).
- Niemi RG, Grofman B, Carlucci C, Hofeller T (1990) Measuring compactness and the role of a compactness standard in a test for partisan and racial gerrymandering. *The Journal of Politics* 52(4):1155–1181.
- Perry WL, McInnis B, Price CC, Smith SC, Hollywood JS (2013) *Predictive policing: The role of crime forecasting in law enforcement operations* (Rand Corporation, Santa Monica, CA).
- Piyadasun T, Kalansuriya B, Gangananda M, Malshan M, Bandara HD, Marru S (2017) Rationalizing police patrol beats using heuristic-based clustering. *2017 Moratuwa Engineering Research Conference (MERCon)*, 431–436 (IEEE).
- Ríos-Mercado RZ (2020) *Optimal Districting and Territory Design* (Springer, Switzerland).
- Rosen S (1974) Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy* 82(1):34–55.
- Sarac A, Batta R, Bhadury J, Rump C (1999) Reconfiguring police reporting districts in the city of buffalo. *OR Insight* 12(3):16–24.
- Saunders J, Hunt P, Hollywood JS (2016) Predictions put into practice: a quasi-experimental evaluation of chicago predictive policing pilot. *Journal of Experimental Criminology* 12(3):347–371.
- Shirabe T (2009) Districting modeling with exact contiguity constraints. *Environment and Planning B: Planning and Design* 36(6):1053–1066.
- United States Census Bureau (2019) 2019 population estimates by age, sex, race, and hispanic origin. <https://www.census.gov/newsroom/press-kits/2019/detailed-estimates.html>.
- Vickrey W (1961) On the prevention of gerrymandering. *Political Science Quarterly* 76(1):105–110.
- Yong HP (1988) Measuring the compactness of legislative districts. *Legislative Studies Quarterly* 13(1):105–115.

## Appendix A: Key Assumptions on Dispatching Rule

We make the following assumptions on the dispatching rule: (1) The number of patrol units allocated to each zone is equal to the number of beats; for simplicity, we do not consider backup units. (2) In response to each call for police service, the nearest available patrol unit is dispatched to the scene of the request; a major incident may require multiple patrol units to respond; we ignore such a possibility because these incidents are rare. (3) If no unit is available in the zone of the request, the call enters a queue with other backlogged calls, which are processed according to the first-come-first-served (FCFS) rule. (4) Upon completion of the service, a patrol unit is either dispatched to a call waiting in the queue, or it immediately returns to its home beat if the queue is empty. In the formulation of the stochastic model below, we assume returning to the home beat is instantaneous.

## Appendix B: Police Patrol and Emergency Response Model

In this section, we present a stochastic model for the police patrol and emergency responses. First we model the patrol units as a queueing system; then we derive expressions for several performance measures. We summarize the notations defined in this section in Table B.1.

**Table B.1 Summary of Notations**

Section	Notation	Description
Setup	$\mathcal{I} = \{i = 1, \dots, I\}$	Set of all beats
	$\mathcal{K} = \{k = 1, \dots, K\}$	Set of all zones
	$D = (d_{ik})_{i \in \mathcal{I}, k \in \mathcal{K}}$	Districting design matrix indicating if beat $i$ is assigned to zone $k$ , $i \in \mathcal{I}$ , $k \in \mathcal{K}$
	$\mathcal{J}^{(k)}$	Set of beats assigned to zone $k \in \mathcal{K}$
	$N^{(k)}$	Total number of beats assigned to zone $k \in \mathcal{K}$
	$\lambda_i$	911 calls-for-service arrival rate in beat $i \in \mathcal{I}$
	$\mu$	Service rate of each response unit
Model	$B^{(k)}, S^{(k)}$	The unsaturated and saturated states for zone $k \in \mathcal{K}$
	$H^{(k)} = (\eta_{nj}^{(k)})$	Optimal unit dispatched to beat $j \in \mathcal{J}^{(k)}$ under unsaturated states $B_n^{(k)}$
	$Q^{(k)} = (q_{nm}^{(k)})$	Transition rate matrix associated with unsaturated states
	$T = (\tau_{ij})$	Average travel time from beat $i$ to beat $j$ , $\forall i, j \in \mathcal{I}$
Metrics	$E_{ij}^{(k)}$	Set of states in which unit $i \in \mathcal{J}^{(k)}$ is an optimal unit to assign to a call from beat $j \in \mathcal{J}^{(k)}$
	$\rho_{ij}^{(k)}$	Fraction of all calls that send unit $i \in \mathcal{J}^{(k)}$ to beat $j \in \mathcal{J}^{(k)}$ in zone $k$
	$\xi_i^{(k)}$	Expected travel time of unit $i \in \mathcal{J}^{(k)}$
	$w^{(k)}(D)$	Total workload of zone $k \in \mathcal{K}$ given districting $D$

## Model Formulation

*Zone Districting.* Consider the police districting problem with  $K$  zones and  $I$  beats ( $K < I$ ). Let  $k \in \mathcal{K} = \{1, \dots, K\}$  denote the set of zones. Let  $i \in \mathcal{I} = \{1, \dots, I\}$  denote both the set of beats, as well as the set of police patrol units that are assigned to these beats. We assume that the shape of each beat is given. A particular zone districting design is thus a graph partition of  $I$  beats into  $K$  zones. Let the binary decision

variable  $d_{ik} \in \{0, 1\}$  denote whether beat  $i$  is assigned to zone  $k$ . The districting decision is represented by a matrix  $D = (d_{ik}) \in \{0, 1\}^{I \times K}$ . For each  $i \in \mathcal{I}$ , we have  $\sum_{k=1}^K d_{ik} = 1$ . Given a zone design  $D$ , the set of beats assigned to zone  $k$  is denoted by  $\mathcal{J}^{(k)} = \{i \mid d_{ik} = 1\} \subseteq \mathcal{I}$  and the number of beats in zone  $k$  is denoted by  $N^{(k)} = \sum_{i \in \mathcal{I}} d_{ik}$ . Throughout this section, we assume the zone districting matrix  $D$  is fixed.

*Hypercube Queue Model.* For a given zone  $k \in \mathcal{K}$ , recall that the set of patrol units in this zone is denoted by  $\mathcal{J}^{(k)}$ . Let  $N^{(k)} = |\mathcal{J}^{(k)}|$  be the total number of patrol units, which is equal to the number of beats in zone  $k$ . We assume that the arrivals of 911 calls in beat  $i \in \mathcal{I}$  follow a time-homogeneous Poisson process with rate  $\lambda_i$ . Let  $\lambda^{(k)} = \sum_{i \in \mathcal{J}^{(k)}} \lambda_i$  be the aggregate call arrival rate in zone  $k$ . Denote the average travel time between two arbitrary beats as a matrix  $T = (\tau_{ij}) \in \mathbb{R}_+^{I \times I}$ , where  $\tau_{ij}$  is the average travel time from beat  $i$  to beat  $j$ . The service time of each patrol unit is independent and identically distributed and follows an exponential distribution with mean  $1/\mu$ . We require  $\lambda^{(k)} < \mu N^{(k)}$  for the stability of this queueing system.

The system state depends on the status of all the units in this zone, and the number of calls in queue to be processed. Based on whether any unit is available, the state space (denoted by  $C^{(k)}$ ) can be divided into two parts: (1) *Unsaturated states* are states where the queue for unprocessed calls is empty. These states can be represented by a hypercube in dimension  $N^{(k)}$  (hence the name “hypercube queue model”). Each vertex of the hypercube corresponds to a state  $B^{(k)} = (b_i)_{i \in \mathcal{J}^{(k)}}$ , where unit  $i$  is busy processing a call if  $b_i = 1$  and idle if  $b_i = 0$ . Considering the state  $B^{(k)}$  as a vector of binary numbers, we index the state by its corresponding decimal form  $B_n^{(k)}$  where  $n = \sum_{j=1}^{N^{(k)}} 2^{j-1} b_j$ . (2) *Saturated states* are states where the queue for unprocessed calls is nonempty. By definition, all patrol units are busy under saturated states. We denote these states by  $\{S_n^{(k)}\}_{n \geq 1}$ ; the state  $S_n^{(k)}$  indicates that there are exactly  $n$  calls waiting in queue for zone  $k$ .

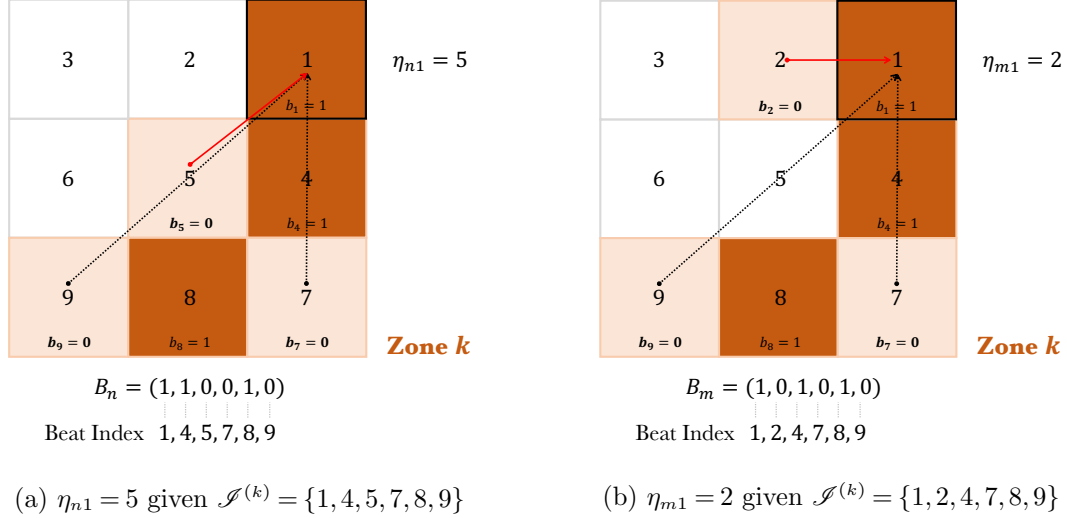
We now define the state transition rates. For the saturated states, the transition rate from  $S_n^{(k)}$  to  $S_{n+1}^{(k)}$ , as well as from the transition rate from  $B_{2^{N^{(k)}}-1}^{(k)}$  to  $S_1^{(k)}$ , is  $\lambda^{(k)}$ . The transition rate from  $S_{n+1}^{(k)}$  to  $S_n^{(k)}$ , as well as the transition rate from  $S_1^{(k)}$  to  $B_{2^{N^{(k)}}-1}^{(k)}$ , is  $\mu N^{(k)}$ .

For the unsaturated states, we define the transition rate matrix  $Q^{(k)} = (q_{nm}^{(k)})$  for zone  $k$ , where  $q_{nm}^{(k)}$  (for  $n \neq m$ ) denotes the rate departing from the  $n$ -th unsaturated state and arriving in the  $m$ -th unsaturated state. Diagonal entries  $q_{nn}^{(k)}$  are defined such that  $q_{nn}^{(k)} = -\sum_{m:m \neq n} q_{nm}^{(k)}$  and therefore the rows of the matrix sum to zero.

There are two classes of transitions on the hypercube: *upward* transitions, which change a unit’s status from idle to busy and *downward* transitions, which change a unit’s status from busy to idle. Let  $d_{nm}$  denote the Hamming distance between two vertices  $B_n$  and  $B_m$  on the hypercube. We define the “upward” Hamming distance as  $d_{nm}^+ = |B_n^C \cap B_m|$  and the “downward” Hamming distance as  $d_{nm}^- = |B_n \cap B_m^C|$ , where  $|B|$  denote the number of binary “ones” in  $B$  (i.e., the number of busy units), and  $B^C$  denote the complement of  $B$ . The downward transition rate from state  $B_n$  to the adjacent states  $B_m$  with  $d_{nm}^- = 1$  is always  $q_{nm}^{(k)} = \mu$ . The upward transition rates, however, will depend on the dispatch rule (denoted by  $H^{(k)}$ ) when a call is received. As we mention in the *Background of Police Operations* section, we assume that when a call is ready to be processed, the nearest available unit with the minimum mean travel time to the location of the call is dispatched. That is, if the call is in beat  $i$  and the system state is  $B_n^{(k)}$ , the index of the unit dispatched is  $\eta_{ni}^{(k)} := \arg \min_{j: b_j=0, b_j \in B_n^{(k)}} \tau_{ji}$  (we assume there are no ties in the mean travel time matrix). The upward

transition rate from state  $B_n$  to an adjacent states  $B_m$  with  $d_{nm}^+ = 1$ , can be obtained by  $q_{nm}^{(k)} = \sum_{i: \eta_{ni}^{(k)} = j, \forall i} \lambda_i$ , where  $j$  is the index of the unit dispatched and satisfies  $m = n + 2^{j-1}$ . To illustrate how this dispatch rule works, we give a detailed example in Figure B.1.

**Figure B.1 Two Examples Show How Districting Decision Affects the Dispatch Rule**



*Notes.* Each numbered block represents a beat. The colored area represents a zone comprised of six beats. The darker blocks represent the corresponding patrol unit  $i$  is busy ( $b_i = 1$ ). Euclidean distance is used as a proxy for travel time between beats. The arrows indicate all possible dispatches, and the red line indicates the optimal dispatch with the shortest travel time.

The steady state probabilities of this system are determined from the balance equations. Let  $\mathbb{P}\{B_m^{(k)}\}$  and  $\mathbb{P}\{S_n^{(k)}\}$  denote the probability that system is occupying state  $B_m^{(k)}$  and  $S_n^{(k)}$  under steady state conditions. For unsaturated states, we have

$$\begin{aligned} (\lambda^{(k)} + \mu |B_m^{(k)}|) \mathbb{P}\{B_m^{(k)}\} &= \sum_{\{B_n^{(k)}: d_{nm}^+ = 1\}} q_{nm} \mathbb{P}\{B_n^{(k)}\} + \sum_{\{B_n^{(k)}: d_{nm}^- = 1\}} \mu \mathbb{P}\{B_n^{(k)}\}, \quad \forall m = 1, \dots, 2^{N^{(k)}} - 2, \\ (\lambda^{(k)} + \mu N^{(k)}) \mathbb{P}\{B_m^{(k)}\} &= \sum_{\{B_n^{(k)}: d_{nm}^+ = 1\}} q_{nm} \mathbb{P}\{B_n^{(k)}\} + \mu N^{(k)} \mathbb{P}\{S_1^{(k)}\}, \quad \text{for } m = 2^{N^{(k)}} - 1, \end{aligned} \quad (\text{B.1})$$

where  $\lambda^{(k)} = \sum_{i \in \mathcal{S}^{(k)}} \lambda_i$  is the aggregate call arrival rate in zone  $k$ . For saturated states, we have

$$(\lambda^{(k)} + \mu N^{(k)}) \mathbb{P}\{S_n^{(k)}\} = \lambda^{(k)} \mathbb{P}\{S_{n-1}^{(k)}\} + \mu N^{(k)} \mathbb{P}\{S_{n+1}^{(k)}\}, \quad \forall n \in \mathbb{Z}_+.$$

In addition, we require that the probability sum to one, namely

$$\sum_{m=0}^{2^{N^{(k)}}-1} \mathbb{P}\{B_m^{(k)}\} + \sum_{n=1}^{\infty} \mathbb{P}\{S_n^{(k)}\} = 1.$$

We note that although the number of saturated states is infinite, it can be expressed in closed form as in an M/M/c queue. The main challenge for computing the steady state distribution is that the number of *unsaturated* states on the hypercube grows exponentially with the number of beats. In the *Efficient Computation Method* section, we will propose an efficient computation method to approximate the steady state distribution, which exploits the fact that the hypercube queue model has a sparse transition matrix.

## Measures of Performance

*Fraction of Dispatches.* To understand how often a patrol unit travels to different beats in the zone, it is critical to know how many calls in each beat have been assigned to a given unit. The fraction of dispatches in zone  $k$  that send unit  $i \in \mathcal{J}^{(k)}$  to beat  $j \in \mathcal{J}^{(k)}$ , denoted by  $\rho_{ij}^{(k)}$ , can be divided into two terms: (1)  $\rho_{ij}^{(k,1)}$  the fraction of dispatches that incurs no queue delay, and (2)  $\rho_{ij}^{(k,2)}$  the fraction of dispatches that incurs a positive queue delay. The fraction of dispatches  $\rho_{ij}^{(k)}$  can be expressed as

$$\rho_{ij}^{(k)} = \underbrace{\sum_{B_n^{(k)} \in E_{ij}^{(k)}} \frac{\lambda_j}{\lambda^{(k)}} \mathbb{P}\{B_n^{(k)}\}}_{\rho_{ij}^{(k,1)}} + \underbrace{\frac{\lambda_j P'_Q}{\lambda^{(k)} N^{(k)}}}_{\rho_{ij}^{(k,2)}}, \quad (\text{B.2})$$

where  $P'_Q = \sum_{m=1}^{\infty} \mathbb{P}\{S_m^{(k)}\} + \mathbb{P}\{B_{2^{N^{(k)}}-1}^{(k)}\}$  specifies the probability that a randomly arriving call incurs a queue delay; the set  $E_{ij}^{(k)} = \{B_n^{(k)} \mid \eta_{nj}^{(k)} = i, n = 1, \dots, 2^{N^{(k)}} - 1\}$  contains all the unsaturated states in which unit  $i$  is the nearest available unit to assign to a call from beat  $j$ , which is determined by the dispatch rule  $H^{(k)} = (\eta_{ni}^{(k)})$ . Equation (B.2) can be found in Larson (1974a), so we omit the proof. Note that in Equation (B.2), the number of beats in zone  $k$  ( $N^{(k)}$ ), the dispatch rule ( $H^{(k)}$ ), and the steady state probability ( $\mathbb{P}\{B_n^{(k)}\}$ ,  $\mathbb{P}\{S_m^{(k)}\}$ ) depends on the districting decision.

*Expected Travel Time.* Another key measure of performance is the travel time that a unit takes for a single dispatch. The travel time is affected by both the beat-to-beat travel time  $\tau_{ij}$  and the fraction of dispatches  $\rho_{ij}^{(k)}$ . The beat-to-beat time  $T = (\tau_{ij})$  can be estimated from the real traffic data, which we discuss below. The fraction of dispatches critically depends on the utilization factor of the hypercube queue system. If the utilization is low, most calls can be processed by units close to their home beats, resulting in short travel time. However, if the system is congested, calls must wait in the queue, which will be processed on an FCFS basis. Each patrol unit is equally likely to process calls that incur a queue delay, regardless of their locations. This will lead to longer travel time.

To derive the equation for the expected travel time for zone  $k \in \mathcal{K}$ , we denote the expected travel time of unit  $i$  for a single dispatch as  $\xi_i^{(k)}$ , the expected travel time for a call in beat  $j$  as  $T_j^{(k)}$ , and the zone-wide mean travel time as  $\bar{T}^{(k)}$ . Using the result in the fraction of dispatches,  $\xi_i^{(k)}$  can be expressed as the weighted sum of mean travel times  $\tau_{ij}$  for all beats  $j$  in zone  $k$  given the dispatch fraction  $\rho_{ij}^{(k)}$ . Specifically, as Larson (1974a) and Larson and Odoni (1981) discuss, the average travel time of unit  $i$  in zone  $k$  can be expressed as

$$\xi_i^{(k)} = \frac{\sum_{j \in \mathcal{J}^{(k)}} \rho_{ij}^{(k,1)} \tau_{ij} + \bar{T}_Q^{(k)} P'_Q / N^{(k)}}{\sum_{j \in \mathcal{J}^{(k)}} \rho_{ij}^{(k,1)} + P'_Q / N^{(k)}}, \quad (\text{B.3})$$

where  $\bar{T}_Q^{(k)} = \sum_{i \in \mathcal{J}^{(k)}} \sum_{j \in \mathcal{J}^{(k)}} \lambda_i \lambda_j \tau_{ij} / (\lambda^{(k)})^2$  is the mean travel time for calls that incur a positive queue delay and  $P'_Q$  is defined in Equation (B.2). Similarly, the expected travel time for calls in beat  $j$  is

$$T_j^{(k)} = \frac{\sum_{i=1}^{N^{(k)}} \rho_{ij}^{(k,1)} \tau_{ij}}{\sum_{i=1}^{N^{(k)}} \rho_{ij}^{(k,1)}} (1 - P'_Q) + \sum_{i=1}^{N^{(k)}} \frac{\lambda_i \tau_{ij}}{\lambda^{(k)}} P'_Q.$$

The zone-wide mean travel time is

$$\bar{T}^{(k)} = \sum_{i=1}^{N^{(k)}} \sum_{j \in \mathcal{J}^{(k)}} \rho_{ij}^{(k,1)} \tau_{ij} + P'_Q \bar{T}_Q^{(k)}.$$

*Response Time.* The response time for a call includes both the travel time for the patrol unit to arrive at the scene and the waiting time if the call incurs a positive queue delay. Therefore, the average response time for calls in beat  $j$  is

$$T_j^{(k)} + \sum_{m=1}^{\infty} \frac{m+1}{\mu N^{(k)}} \mathbb{P}\{S_m^{(k)}\} + \frac{1}{\mu N^{(k)}} \mathbb{P}\{B_{2^{N^{(k)}}-1}^{(k)}\}.$$

The infinite sum in the second term can be solved in closed form as in an M/M/c queue. The zone-wide mean response time is

$$\bar{T}^{(k)} + \sum_{m=1}^{\infty} \frac{m+1}{\mu N^{(k)}} \mathbb{P}\{S_m^{(k)}\} + \frac{1}{\mu N^{(k)}} \mathbb{P}\{B_{2^{N^{(k)}}-1}^{(k)}\}.$$

*Patrol Unit Workload.* Finally, we consider the workload of the patrol units in each zone. In the following optimization model, we will consider districting designs that balance the workload across zones; therefore, this metric will play a vital role in defining the optimization objective. The workload of patrol unit  $i$  in zone  $k$  is

$$w_i^{(k)} = \sum_{n: (B_n^{(k)})_i=1} \mathbb{P}\{B_n^{(k)}\} + \sum_{m=1}^{\infty} \mathbb{P}\{S_m^{(k)}\}.$$

Here, we introduce the districting decision  $D$  in defining the zone workload and will use it in the objective function of our redistricting optimization. The total workload for all patrol units in zone  $k$  is simply

$$w^{(k)}(D) = \sum_{i \in \mathcal{J}^{(k)}} w_i^{(k)}. \quad (\text{B.4})$$

Note that the workload for zone  $k$  depends on the decision variable  $D$  through  $\mathcal{J}^{(k)}$ . A detailed explanation is given below.


*Dependence of Performance Measures on the Districting Decision.* A zone districting decision  $D$  divides the entire geographical region into  $K$  police zones, or  $K$  hypercube queue models. The dynamics between the above performance measures and the beat-to-zone allocations are complex, because the parameters of  $K$  hypercube queues are endogenously determined by the decision of beat allocation. As Figure B.2 shows, the police zone workload  $w^{(k)}(D)$  depends on decision  $D$  through a long dependence chain. Specifically, given a particular decision  $D$ , the beats will be divided into  $K$  zones, where zone  $k \in \mathcal{K}$  corresponds to a set of beats  $\mathcal{J}^{(k)}(D)$ . This leads to a unique combination of  $K$  hypercube queues, where their state spaces  $\{C^{(k)}\}_{k \in \mathcal{K}}$  and dispatch preferences  $\{H^{(k)}\}_{k \in \mathcal{K}}$  are jointly controlled by how decision  $D$  partitions the beats. By solving the linear equation system formed by the generator matrix  $Q^{(k)}$  for each hypercube, we can obtain  $K$  steady state distributions  $\{\pi^{(k)}\}_{k \in \mathcal{K}}$ . Accordingly, we can compute the performance measures  $\xi_i^{(k)}(D), \rho_{ij}^{(k)}(D)$  for each hypercube queue based on their steady state distribution. The unknown police zone workload  $w^{(k)}(D)$  for  $K$  zones given a particular decision can be evaluated through the above process.

### Efficient Computation Method

Performing workload estimation efficiently is critical in searching for the optimal zone design in the solution space. Take Atlanta as an example; there are more than 2,000 possible ways to shift the existing design even with less than only the four changes that are allowable. Each zone in one particular design corresponds to a unique hypercube queue, where its zone workload needs to be re-evaluated separately.



**Figure B.2** The Dependence Diagram for the Key Variables in the  $K$ -Hypercube Model



A critical step for obtaining police zone workload estimation given a decision is to compute the steady state distribution. Let  $\lambda^{(k)} = \sum_{i \in \mathcal{J}^{(k)}} \lambda_i$  be the total call arrival rate in zone  $k$ . It is clear that for zone  $k$ , if the system is in some state  $B_n^{(k)}$  with  $|B_n^{(k)}|$  servers busy, the total transition rate “downward” to states with one less busy server is  $|B_n^{(k)}|\mu$  and the total transition rate “upward” to states with one more busy server is  $\lambda$  (except for  $|B_n^{(k)}| = N^{(k)}$ , in which case the upward transition rate is zero). Thus, if we aggregate states according to the numbers of servers busy, we obtain the well-known M/M/c multiserver queuing system. Then the number of steady state probabilities in the saturated state set  $\{S_n^{(k)}\}_{n \geq 1}$  is given by closed-form expression:

$$\mathbb{P}\{S_n^{(k)}\} = \frac{\frac{(\lambda^{(k)}/\mu)^{N^{(k)}}}{N^{(k)}!} \left(\frac{\lambda^{(k)}}{\mu N^{(k)}}\right)^n}{\sum_{s=0}^{N^{(k)}-1} \frac{(\lambda^{(k)}/\mu)^s}{s!} + \frac{(\lambda^{(k)}/\mu)^{N^{(k)}}}{N^{(k)}!} \frac{1}{1 - \lambda^{(k)}/(\mu N^{(k)})}}, \quad \text{for all } n \in \mathbb{Z}_+.$$

Let the row vector  $\pi^{(k)} = (\mathbb{P}\{B_m^{(k)}\})$  denote the steady state probabilities in the unsaturated state set  $C^{(k)}$ . Then the balance equations (B.1) can be written compactly as

$$\pi^{(k)} = \pi^{(k)}(I^{(k)} + Q^{(k)}/\gamma^{(k)}),$$

where  $I^{(k)}$  is the identity matrix (of the same size as  $Q^{(k)}$ ) and  $\gamma^{(k)} = \max |q_{ii}^{(k)}|$  is the largest absolute value on the diagonal of  $Q^{(k)}$ . The matrix  $(I^{(k)} + Q^{(k)}/\gamma^{(k)})$  is the probability transition matrix of a discrete-time Markov chain obtained through uniformization of the original continuous-time Markov chain.

Theoretically, the solution to this set of equations requires a matrix inversion. However, the size of the matrix is equal to  $2^{2N^{(k)}}$ , thus becoming huge for even moderate values of  $N^{(k)}$ . Fortunately, due to the sparsity of the transition rate matrix, we can find the stationary distribution  $\pi^{(k)}$  iteratively by the power method. More specifically, given the transition rate matrix  $Q^{(k)}$  and an initial vector  $\pi_0^{(k)}$ , compute  $\pi_t^{(k)} = \pi_{t-1}^{(k)}(I^{(k)} + Q^{(k)}/\gamma^{(k)})$  for  $t = 1, 2, \dots$ , until the distance (e.g., the sup norm) between  $\pi_t^{(k)}$  and  $\pi_{t-1}^{(k)}$  is small enough.

## Appendix C: Model Estimation and Workload Prediction

We obtained the three data sets described above for five years from 2013 to 2017. This enabled us to fit the hypercube queue model to predict the future workload.

*Call Arrival Rate Prediction.* Predicting the call arrival rate is particularly challenging. Although our operations model assumes the call arrival rates are time homogeneous, we observe in the actual data that the

call arrival rates have a significant seasonality pattern and yearly trend, as we show in Figure 7 (b), and a correlation over adjacent geographical areas. Therefore, we propose a spatio-temporal model to predict future call arrival rates.

We index the five years from 2013 to 2017 by  $\ell \in \mathcal{L} = \{1, \dots, L\}$ ,  $L = 5$ . Let  $\lambda_{i\ell} \in \mathbb{R}_+$  represent 911 call arrival rates of beat  $i$  in year  $\ell$ . Let vector  $\mathbf{c}_{i\ell} \in \mathbb{R}^M$  represent the values of  $M$  demographic factors of beat  $i$  in year  $\ell$ . The graph  $\mathcal{G}$  is given by associating a node with every beat and connecting two nodes by an edge whenever the corresponding beats are geographically adjacent. The set of adjacency pairs is defined by  $\mathcal{A} = \{(i, j) \in \mathcal{I} : i, j \text{ are adjacent in } \mathcal{G}\}$ . Then, we use the spatially lagged endogenous regressors (Rosen 1974) defined as

$$\lambda_{i\ell} = \underbrace{\sum_{(i,j) \in \mathcal{A}} \alpha_{ij} \lambda_{j\ell}}_{(\dagger)} + \underbrace{\beta_0 \lambda_{i,\ell-1}}_{(\ddagger)} + \underbrace{\sum_{t=0}^p \mathbf{c}_{i,\ell-t} \boldsymbol{\beta}_t}_{(\ddagger\ddagger)} + \epsilon_i, \quad \ell = p+1, \dots, L, \quad (\text{C.1})$$

where  $(\dagger)$  represents the influence of neighboring beats  $j : (i, j) \in \mathcal{A}$ ,  $(\ddagger)$  represents the influence of the arrival rate in the last year, and  $(\ddagger\ddagger)$  represents the influence of demographic factors in the past  $p$  years. The  $p$  is the total number of past years of data that we consider for fitting the regressor (in our case,  $p = 1$ ). The spatial coefficient matrix  $A = (\alpha_{ij}) \in \mathbb{R}^{I \times I}$  specifies the spatial correlation between arrival rates of two beats. The temporal coefficient  $\beta_0 \in \mathbb{R}$  specifies the influence of the arrival rate in the last year. The coefficient  $\boldsymbol{\beta}_t \in \mathbb{R}^M, \forall 1 \leq t \leq p$  specifies the correlation of demographic factors. Let  $X_i$  denote a subset of the data  $\{\{\mathbf{c}_{i\ell}\}_{\ell \in \mathcal{L}}, \{\lambda_{j\ell}\}_{(i,j) \in \mathcal{A}, \ell \in \mathcal{L}}\}$ . The error  $\epsilon_i \in \mathbb{R}$ ,  $i \in \mathcal{I}$  is a random noise,  $\mathbb{E}[\epsilon_i | X_i] = 0$  and  $\text{Cov}[\epsilon_i, \epsilon_j | X_i, X_j] = \Sigma_\theta(i, j)$  for all  $i, j \in \mathcal{I}$ . The spatial covariance of the noise between two beats  $i, j$  is determined by a correlation function  $\Sigma_\theta$ , which is a function of their spatial distance  $s_{ij}$ , and is parameterized by  $\theta$ . Some commonly used spatial models, including the Gaussian model (Lee et al. 2014), exponential model (Gaetan and Guyon 2010), and Matérn model (Gaetan and Guyon 2010). Here we adopt the exponential model where  $\Sigma_\theta(i, j) = \theta_1 \exp\{-\theta s_{ij}\}$ , where  $\theta$  is a prespecified parameter. The other parameters of the model including  $\{A, \beta_0, \{\boldsymbol{\beta}_t\}_{1 \leq t \leq p}\}$  can be fitted by maximum likelihood estimation using historical data  $\{X_i\}_{i \in \mathcal{I}}$ .

*On-scene Time Estimation.* Because we assume the on-scene time of a response unit is independent of its service region, we investigate the distribution of the on-scene time of each unit in the same zone. By looking into the data, we observe that the on-scene time of response units follows an exponential distribution shown in Figure 5. We obtain the mean on-scene time  $1/\mu = 31.2$  (minutes) by averaging the on-scene time of all dispatches in the past.

*Travel Time Estimation.* We estimated the beat-to-beat travel time  $T = (\tau_{ij})$  from the 911 call dispatch data. The estimation result is shown (Figure 6) by averaging the travel time of all 911 call dispatches that started in beat  $i$  and ended in beat  $j$ . In this figure, beats are ordered according to their relative location in the city. Thus, the beats in the same zone are consecutively numbered. Clearly, the figure reveals that local traveling within a neighborhood region usually takes a longer time than traveling across zones. Two major reasons account for this phenomenon. First, cross-zone calls usually have higher priority, which takes less time for a response and requires more patrol units to assist the call; therefore, they sometimes need police

from other zones. Second, response units frequently use Atlanta's highway system when responding to the non local calls, thus reducing the travel time.

## Appendix D: Literature Review on Police Workload

In our study, we consider balancing police workload as one of the design objectives. Various definitions for police workload have been adopted in the literature. Mitchell (1972) defines police workload as the sum of service time and travel time. Curtin et al. (2005) use the number or frequency of 911 calls occurring at each district as a proxy for the workload. In Bodily (1978) and D'Amico et al. (2002), the workload is defined as the fraction of working time that an agent spends attending to calls. Benveniste (1985) considers an equivalent measure. Response time is also an important performance measure representing the time between the arrival of a call for service and the arrival of a unit at the incident location. According to Bodily (1978), the reduction in response time results in several beneficial effects, including (1) an increased likelihood of intercepting a crime in progress, (2) a deterrent effect on criminals, and (3) an increased confidence in the police by citizens.

## Appendix E: Zone Design Optimization

In this section, we develop an optimization model for zone redesign. We assume that all the beat areas are given; therefore, the zone districting decision is equivalent to a graph partition in which we allocate the beats to a fixed number of zones. Recall that the matrix  $D = (d_{ik})_{I \times K}$  represents the beat allocation decisions, and  $w^{(k)}(D)$  represents the mean police workload in zone  $k \in \mathcal{K}$  given a districting design  $D$ .

The goal is to minimize the workload variance subject to some shape constraints (e.g., contiguity and compactness) for each zone. The zone redesign problem can be expressed as

$$\begin{aligned} & \underset{D \in \{0,1\}^{I \times K}}{\text{minimize}} & f(D) &:= \sum_{k=1}^K \left( w^{(k)}(D) - \frac{1}{K} \sum_{k'=1}^K w^{(k')}(D) \right)^2 \\ & \text{subject to} & \sum_{k=1}^K d_{ik} &= 1, \quad \forall i \in \mathcal{I}, \\ & & & \text{contiguity and compactness constraints for zone } k \in \mathcal{K}. \end{aligned} \tag{E.1}$$

Notice that the zone-level workload variance is chosen as the objective function in problem (E.1), because APD suggested this metric for its zone redesign. However, one may also consider other performance metrics, such as mean response time or workload variance at the beat level (we provided expressions for these metrics in Appendix B). Our framework can be similarly adapted to solve these problems.

### Zone Workload Approximation

We consider a local approximation of the workload function, which allows the objective function of the problem (E.1) to be expressed explicitly. Consider a given zone with an existing design. If we make a small change to this zone (e.g., by adding a few beats or removing a few beats), the workload of this zone can be approximated by its first-order Taylor expansion as

$$w^{(k)}(D) = \theta_{0k} + \sum_{i \in \mathcal{I}} \theta_{ik} d_{ik} + \tilde{\delta}^{(k)}, \quad \forall k \in \mathcal{K} \tag{E.2}$$

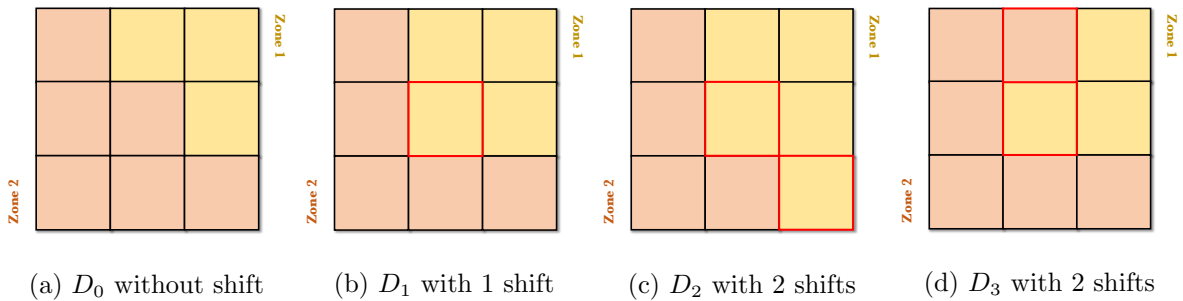
where  $\theta = \{\{\theta_{0k} \in \mathbb{R}\}_{k \in \mathcal{K}}, \{\theta_{ik} \in \mathbb{R}\}_{i \in \mathcal{I}, k \in \mathcal{K}}\}$  denotes the coefficients of this approximation model, and  $\tilde{\delta}^{(k)}, \forall k \in \mathcal{K}$  represent the approximation errors.

We estimate the parameters in Equation (E.2) as follows. First, for each of the six zones in Atlanta, we sample 1,000 designs, which we created using small perturbations from the existing design. In particular, these samples are generated by randomly changing a few beats on the boundary of the existing zone design (Figure E.1). Obtaining a single sample involves the following processes: (1) generate a random design by changing a limited number of beat assignments; (2) compute the steady state probabilities of six zones by solving Equation (B.1); (3) obtain simulated zone workload by computing Equation (B.4). The computational bottleneck of getting one sample lies in Step (2), which might take two minutes or so, depending on the number of the beat shifts and the computer's processing power. The entire generating process takes approximately 33 hours. Next, we estimate the parameters of the linear function in Equation (E.2) using generated data samples. More specifically, we apply the least squares method to obtain the estimated values  $\hat{\theta}$ , where the residual is  $R^2 = 0.987$ . Finally, we replace the exact workload function  $w^{(k)}(D)$  with its approximation  $\hat{w}^{(k)}(D|\hat{\theta}) = \hat{\theta}_{0k} + \sum_{i \in \mathcal{I}} \hat{\theta}_{ik} d_{ik}$ . The objective function  $f(D)$  of the optimization problem (E.1) is thus approximated by

$$\tilde{f}(D|\hat{\theta}) = \sum_{k=1}^K \left( \hat{w}^{(k)}(D|\hat{\theta}) - \frac{1}{K} \sum_{k'=1}^K \hat{w}^{(k')}(D|\hat{\theta}) \right)^2. \quad (\text{E.3})$$

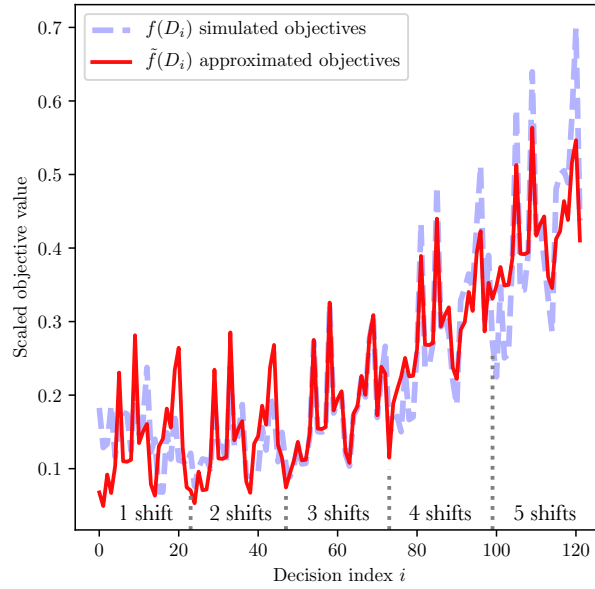
In Figure E.2, we evaluate the effectiveness of our approximation by plotting the values of the objective function (i.e., workload variance) and corresponding approximations for 122 randomly selected *out-of-sample* designs. These designs are also created based on the existing design in Atlanta. Of these designs, 23 of them have one beat shift, 24 have two beats shifts, 26 have three and four beats shifts, and the others have five beat shifts. Note that these designs are uniformly sampled from all possibilities and not necessarily optimal for the objective. The blue dash line plots the exact objective value  $\{f(D_i)\}$ , calculated by simulation. The red curve plots the approximated workload value  $\{\tilde{f}(D_i|\hat{\theta})\}$ . The result shows that the approximated objective values match well with the true objective values for these randomly generated designs.

**Figure E.1 Examples Show Three Local Decisions by Making 1, 2, 2 Shifts, Respectively**



*Notes.* Each square block represents a single beat. Blocks with the same color are in the same zone. The red boxes indicate the shifts based on the original decision. (E.1a) shows the original decision  $D_0$  without any shift; (E.1b) shows the decision  $D_1$  with only one shift on the boundary; (E.1c, E.1d) show two decisions  $D_2, D_3$  with different shifts on the boundary.

**Figure E.2** The Approximated and Simulated Objective Values Versus a Subset of  $\{D_i\}$ , Which Differs to Each Other by Shuffling a Small Number of Beats



*Notes.* Decisions indexed by  $0 \leq i < 23$ ,  $23 \leq i < 47$ ,  $47 \leq i < 73$ ,  $73 \leq i < 99$ ,  $99 \leq i < 122$  correspond to randomly generated designs with 1, 2, 3, 4, and 5 shifts, respectively, comparing to the existing zone plan regarding the beat assignment.

### Linearizing the Objective Function

With the parametric approximation (E.3), the objective function measuring the workload variance becomes a *quadratic* function of the binary decision variables  $D$ . Using a standard technique for quadratic binary variables, this quadratic function can be expressed as a linear form by introducing auxiliary variables  $e_{ijkk'}$  with additional constraints

$$e_{ijkk'} \leq d_{ik}, e_{ijkk'} \leq d_{jk'}, e_{ijkk'} \geq d_{ik} + d_{jk'} - 1, e_{ijkk'} \geq 0 \quad \text{for all } i, j \in \mathcal{I}, k, k' \in \mathcal{K}.$$

These constraints ensure that  $e_{ijkk'} = d_{ik}d_{jk'}$ . Therefore, the objective function can be written as

$$\begin{aligned} \tilde{f}(D|\hat{\theta}) = & \sum_{k=1}^K \left[ c^{(k)} + \sum_{i=1}^I \sum_{j=1}^I \hat{\theta}_{ik} \hat{\theta}_{jk} e_{ijkk} + 2\hat{\theta}_{0k} \sum_{i=1}^I \hat{\theta}_{ik} d_{ik} - \right. \\ & \frac{2}{K} \left( \sum_{k'=1}^K \hat{\theta}_{0k'} \right) \sum_{i=1}^I \hat{\theta}_{ik} d_{ik} - \frac{2}{K} \hat{\theta}_{0k} \sum_{k'=1}^K \sum_{i=1}^I \hat{\theta}_{ik'} d_{ik'} - \frac{2}{K} \sum_{i=1}^I \sum_{j=1}^I \sum_{k'=1}^K \hat{\theta}_{ik} \hat{\theta}_{jk'} e_{ijkk'} + \\ & \left. \frac{2}{K^2} \left( \sum_{k'=1}^K \hat{\theta}_{0k'} \right) \sum_{k'=1}^K \sum_{i=1}^I \hat{\theta}_{ik'} d_{ik'} + \sum_{k'=1}^K \sum_{k''=1}^K \sum_{i=1}^I \sum_{j=1}^I \hat{\theta}_{ik'} \hat{\theta}_{jk''} e_{ijk'k''} \right], \end{aligned} \quad (\text{E.4})$$

where  $c^{(k)} = \hat{\theta}_{0k}^2 + 2\hat{\theta}_{0k} \sum_{k'=1}^K \hat{\theta}_{0k'}/K + (\sum_{k'=1}^K \hat{\theta}_{0k'})^2/K^2$  is a constant term for all  $k \in \mathcal{K}$ . Thanks to the linearization, we can take advantage of the existing optimization techniques (e.g., branch and cut algorithm) to solve the problem with a relatively smaller size to optimality or evaluate the quality of heuristic search solutions.

However, in this work, the reformulated problem contains more than 240,000 variables and is characterized by a large search space with more than  $2.4 \times 10^{59}$  possible solutions. To address the computational complexity of solving this large-scale MILP, consider the implementation feasibility, and avoid drastic changes from the existing zone design, we use a fast local search method based on simulated annealing. The optimization process takes about 30 to 50 minutes depending on where the problem converged given a single laptop's computational power with quad-core processors with speeds up to 4.7 GHz.

### Contiguity and Compactness Constraints

*Contiguity Constraints.* We formulate the contiguity constraints using a network flow approach (see Shirabe 2009). We define a flow network, where each beat is represented by a node. An arc from node  $i$  to node  $j$  exists if the beats associated with these two nodes are geographically adjacent. Let  $\mathcal{A}$  represent the set of arcs.

We now define flow constraints for each zone  $k \in \mathcal{K}$  to ensure contiguity. Let  $n_{\max}$  be the maximum number of beats allowed in a zone. If beat  $i$  belongs to zone  $k$ , i.e.,  $d_{ik} = 1$ , we send one unit of flow into node  $i$ . Let  $v_{ijk}$  be the amount of flow from beat  $i$  to beat  $j$  if  $(i, j) \in \mathcal{A}$  and both beats belong to zone  $k$ . In addition, for each zone  $k$ , we specify a special “sink node.” Let  $h_{ik}$  be 1 if beat  $i \in \mathcal{I}$  is selected as the sink node in zone  $k \in \mathcal{K}$ ; otherwise, let  $h_{ik}$  be 0. The sink node receives one unit of flow from every other node in the same zone. This ensures that there is a path of positive flow from any node in the zone to the sink node, implying contiguity. In sum, we formulate the contiguity constraints as

$$\sum_{j: (i,j) \in \mathcal{A}} v_{ijk} - \sum_{j': (j',i) \in \mathcal{A}} v_{j'ik} \geq d_{ik} - n_{\max} h_{ik}, \quad \forall i \in \mathcal{I}, k \in \mathcal{K}, \quad (\text{E.5a})$$

$$v_{ijk} + v_{jik} \leq (n_{\max} - 1) d_{ik}, \quad \forall (i, j) \in \mathcal{A}, k \in \mathcal{K}, \quad (\text{E.5b})$$

$$\sum_{i=1}^N h_{ik} = 1, \quad \forall k \in \mathcal{K}, \quad (\text{E.5c})$$

$$h_{ik} - d_{ik} \leq 0, \quad \forall i \in \mathcal{I}, k \in \mathcal{K} \quad (\text{E.5d})$$

$$d_{ik}, h_{ik} \in \{0, 1\}, \quad \forall i \in \mathcal{I}, k \in \mathcal{K}, \quad (\text{E.5e})$$

$$v_{ijk} \geq 0, \quad \forall (i, j) \in \mathcal{A}, k \in \mathcal{K}. \quad (\text{E.5f})$$

Specifically, Constraints (E.5a) represent the net outflow from all nodes. The two terms on the left indicate, respectively, the total outflow and total inflow of beat  $i$ . If beat  $i$  is included in zone  $k$  and it is not a sink, then we have  $d_{ik} = 1$ ,  $h_{ik} = 0$ , and thus beat  $i$  must have an outflow of one unit. If beat  $i$  is included in zone  $k$  and is a sink, then we have  $d_{ik} = 1$ ,  $h_{ik} = 1$ , and thus beat  $i$  has a negative outflow lower bounded by  $1 - n_{\max}$ . Constraints (E.5b) ensure that if beat  $i$  is not in zone  $k$  (i.e.,  $d_{ik} = 0$ ), there must be no flow into or out of beat  $i$  with any nodes in zone  $k$ ; if beat  $i$  is in zone  $k$  (i.e.,  $d_{ik} = 1$ ), the flow between two beats does not exceed  $n_{\max} - 1$ . Constraints (E.5c) ensure that each zone must have a unique sink node. Constraints (E.5d) ensure that beat  $i$  must belong to zone  $k$  if it is the sink for that zone. These constraints together ensure that each zone is contiguous.

*Compactness Constraints.* We consider two ways to enforce the compactness of each zone: distance compactness (Niemi et al. 1990, Yong 1988) and shape compactness (Garfinkel and Nemhauser 1970). For distance

compactness, we require the distance between any two beats in a zone to be bounded by some parameters. For shape compactness, we require the square of the diameter divided by the zone's area to be upper bounded by another parameter. These constraints can be formulated as follows.

For each beat  $i \in \mathcal{I}$ , let  $A_i$  be its area. For each pair of beats  $(i, j) \in \mathcal{I} \times \mathcal{I}$ , let  $l_{ij}$  be the square of the distance between the center of these beats. We choose two parameters  $\zeta_1, \zeta_2 > 0$  to control the level of distance compactness and shape constraints, respectively. Then we have

$$l_{ij}e_{ijkk} \leq \zeta_1, \quad \forall i \in \mathcal{I}, j \in \mathcal{I}, k \in \mathcal{K}, \quad (\text{E.6a})$$

$$l_{ij}e_{ijkk} \leq \zeta_2 \sum_{i'=1}^K d_{i'k} A_i, \quad \forall i \in \mathcal{I}, j \in \mathcal{I}, k \in \mathcal{K}. \quad (\text{E.6b})$$

Recall that  $e_{ijkk} = d_{ik}d_{jk}$  is equal to 1 if both beat  $i$  and beat  $j$  are in zone  $k$ , and is equal to 0 otherwise. Constraints (E.6a) represent the distance compactness criterion. Constraints (E.6b) represent the shape compactness criterion, where  $\sum_{i=1}^K d_{ik} A_i$  is the area of zone  $k \in \mathcal{K}$ .