RESEARCH   **Open Access**

Police patrol districting method and simulation evaluation using agent-based model & GIS

Yue Zhang[*] and Donald E Brown

Abstract

Police patrols play an important role in public safety. The patrol district design is an important factor affecting the patrol performances, such as average response time and workload variation. The redistricting or redrawing police command boundaries can be described as partitioning a police jurisdiction into command districts with the constraints such as contiguity and compactness. The size of the possible sample space is large and the corresponding graph-partitioning problem is NP-complete. In our approach, the patrol districting plans generated by a parameterized redistricting procedure are evaluated using an agent-based simulation model we implemented in Java Repast in a geographic information system (GIS) environment. The relationship between districting parameters and response variables is studied and better districting plans can be generated. After in-depth evaluations of these plans, we perform a Pareto analysis of the outputs from the simulation to find the non-dominated set of plans on each of the objectives. This paper also includes a case study for the police department of Charlottesville, VA, USA. Simulation results show that patrol performance can be improved compared with the current districting solution.

Introduction

Police patrols play an important role in public service by responding to incidents, deterring and preventing crimes. It can give a sense of security to people who need protection and discourage those who may commit crimes in the absence of a patrol [1]. Police patrolling is an indispensable component and function of police departments [2]. The goals and objectives of police patrol include crime prevention, criminal apprehension, law enforcement, order maintenance, public services, and traffic enforcement [3]. However, since police resources are limited, there is an understandable interest in patrol strategies and operations that provide safety at minimum cost.

Typically, a city is partitioned into command districts or precincts. Each command district usually has a headquarters and a commanding officer to manage and supervise the police patrol operations. Each command district is further divided into several beats or sectors [4]. Police can effectively manage their operations through the design of the command districts and the choice of patrol strategies of the police units within those districts. There are three major types of patrol strategies for patrol officers: active

patrol, random patrol, and directed patrol. In active patrol, patrol officers should use every opportunity to discover, detect, observe, and interdict the unusual event. Random patrol means patrol routes should be random and varied so that the patrol behaviors will not be predicted by potential criminals. In directed patrol, patrol officers pay more efforts to hotspots of crime so they can respond quickly and reduce the crimes in hotspots. In practice, patrol officers may choose one strategy or combine them to accommodate the specific conditions in their area [3]. Another important patrol principle is beat integrity. Patrol officer is expected to remain within assigned patrol district. Beat integrity can be absolute or relative, depending on the number of patrol units assigned and the size of the patrol district and the activity within the area [3]. For absolute beat integrity, patrol officers should remain in the assigned patrol district at all times. In relative beat integrity, patrol officers remain largely within the patrol area and leave it only for good reason, such as back up another officer or respond to a call for service (CFS) incident in another patrol district [3]. When there is a CFS incident, many police departments use the principle that "the nearest police car responds to the call", which belongs to the case of relative beat integrity.

*Correspondence: yz5yf@virginia.edu
Predictive Technology Laboratory, Department of Systems and Information Engineering, University of Virginia, Charlottesville, Virginia, USA

Police patrol district design or region design is one of the important resource allocation strategies of police department. The deployment of patrol force through designated patrol districts is a standard management method to enhance the deterrent capability of uniformed patrol force [3]. Better districting plans can lead to lower response times, officers' familiarization with their assigned area, more efficient utilization of personnel, equalized distribution of workload, uniform police visibility, enhanced officer safety, balanced police response to calls and officer accountability [3]. For small and medium-sized cities, patrol district design determines the patrol boundaries for each patrol unit, for example, police car. For large cities, patrol region design usually starts from higher level of resource allocation. Larger patrol districts for several patrol units are determined first then these districts are further divided for each patrol unit. Patrol districts should be designed in a rational and systematic manner [3]. Several factors need to be considered when designing patrol districts: 1) the size of the area to be patrolled should not exceed the limit that a single officer can cover, depending on types of the area (urban, suburban or rural), 2) natural or man-made barriers, (e.g., rivers, railroad tracks, interstate highways, valleys), 3) workload indicators (e.g., CFS, criminal incidents, traffic patterns), 4) significant area characteristics (e.g., local neighbor hoods, major shopping centers) [3]. Different cities have different CFS and crime pattern. CFS incidents or crime events are more likely to happen in some area than other areas. So, the police patrol resources should be spatially allocated based on the characteristics of cities. The size, shape, demographics, and geography of the patrol district is a major determinant of the effectiveness of patrol operations and is the focus of this paper.

Police patrol district design

As noted previously the design of police patrol districts can significantly impact the effectiveness of police operations. Typically two major questions constitute the essential elements of patrol design: 1) What areas define each district and 2) What resources should be assigned to each area [5]? This paper concerns the first question and clearly the answer to this question determines the expected performance of the police on a number of performance metrics. The next subsection highlights these metrics. Subsection "Operations research and GIS methods" provides background to previous work in this area and the remaining two subsection show the districting problem reduces the graph selection and then how simulation can be employed to help design districts.

Performance measurements and related problems

There are two common performance measures used for police district design: the average response time and the

variation of workload [4]. Quick response to citizen calls for service can 1) improve the chances of catching the offender at the scene or nearby 2) increase the changes of identify and locate witnesses, 3) provide immediate gathering of physical evidence, 4) provide immediate lifesaving first aid, 5) enhances the reputation of police department, 6) creates citizen satisfaction with the police [5,6]. The spatial distribution and allocation of police cars to districts affect both performance measures. In a 1971 study of the New York Police Department more than half of all dispatches were inter-sector dispatches (between districts). Usually, the nearest police car responds to the CFS incident so it may cross the patrol boundaries to respond. The average response time for these inter-sector responses was approximately 40% greater than that for intra-sector responses and the average travel distance was about 53% greater [7]. This large proportion of inter-sector cases indicates that the district design may not be efficient with regard to the first measure, response time.

Our own interviews with members of the Charlottesville, VA Police Department have shown that the response times do not vary much between or within districts. This is because Charlottesville is small city (approximately 7 miles in diameter) with a population of about 40,000. However, these interviews did reveal large variations in the workloads of officers in different districts. This shows a lack of efficiency with respect to the second objective: variation in workload.

Approaches to police district design

Manual methods

Historically, the geographic patrolling boundaries were drawn by hand based on police department's knowledge and experience of the total area and the availability of the police resources [2,8]. Police department also considered the natural boundaries, such as the hills or rivers, the locations of hotspots of crime as well as the administrative boundaries, such as neighborhoods and communities [9]. The limitation of the hand drawn boundaries is that the human is limited in the number of options they can consider and in their formal evaluation of the alternatives. Each alternative represents a different possibility in terms of workload and response times, but assessing these by hand can be very time consuming. The assessment is particularly difficult since both workload and response times are stochastic. This inability to measure the efficiency of the districting alternatives means that manual methods are not appropriate and good computer-based methods are required in order to create district that can positively impact high-level decision making by the police [2,9].

Operations research and GIS methods

According to [9], the first OR application was to use p-Median clustering as in [8] to minimize the total weighted

travel distance to service the expected calls. In 1979, Aly and Litwhiler used an interchange heuristic method to allocate police briefing stations to districts [10]. In 2002, a simulated annealing approach was used by Amico et al., to search for a good partitioning of the police command districts by assessing the average response time and variation of workload of police officers with a simulation called PCAM [4]. GIS methods can be also used in police district design. In 2010 [9] used GIS methods with a maximal covering formulation to determine optimal police patrol areas. This approach produced good alternatives for police districts design. However, the criteria for evaluation were based on static calculation of distances, weights of incidents, and queuing model statistics. Police patrols are highly dynamic activities. It is difficult to use static approaches to find effective answers for most police. Instead, simulation methods (as discussed in Subsection "Evaluation of districting plans using simulation") of police patrol can provide more useful evaluations.

Problem complexity
Computational complexity
A convenient formulation for the police district design problem is as the aggregation of smaller geographic units into larger units that form the districts. This problem has a rich history in mathematics and operations research. This problem reduces to the graph-partition problem with the constraints of contiguity and compactness. One example of graph-partitioning is the political redistricting problem. The census blocks are partitioned into districts such that each district has equal population. There are two constraints on the problem solution: contiguity and compactness. The polygons within each district must be spatially contiguous so that all members of district can connect with one another without leaving the district. Compactness provides for shorter travel times between members of the district. Evaluation of compactness depends on the distance metric used. For Euclidean distance a circle is the most compact shape, while for Manhattan distance a square is most compact. The graph partitioning problem has been shown to be NP-hard [11].

Another formulation for districting uses graph theory. In this formulation, the regions or geographic units as nodes and the adjacency between regions as edges connect the nodes. For each node, there is a weight representing attributes of the region, for example, population. The objective of districting is to maximize or minimize the function of the weights between different districts. For example, we could use the absolute difference as the function and minimize the difference in population between districts. For the police districts we want to minimize the difference in forecasted activity between districts since

this will minimize the difference in travel or response time. We could also minimize the difference in workload between districts. This formulation of the districting problem can be reduced to the problem of "Cut into Connected Components of Bounded Weight", which is also NP-hard [12].

As with all optimization problems we can employ either exact or heuristic methods [11]. Exact methods systematically exam districting plans explicitly or implicitly. Since the problem is NP-hard, we can use explicit methods only for small districting problems that have limited practical applications. Heuristic methods, such as, simulated annealing, genetic algorithms, and stochastic gradient ascent can find local optima for the districting problem. These local optima can provide good operating solutions for police departments. The approach described in this paper in Section "Approach to police patrol districting" is a heuristic.

Evaluation complexity
The police patrol district design problem has greater complexity than the computational issues mentioned in Subsection "Computational complexity". Consider the graph-theoretic formulation of the districting problem. In this formulation some function of the region weights is used as the optimization criterion. For example, the sum of the absolute difference in populations between all pairs of regions can be used. In the police districting problem the evaluation is more complicated than a simple sum of average differences. To evaluate the response time we need to compute the travel times for the incidents within each district that require police attention. The locations and times of these incidents are stochastic and must be modeled by an empirical distributions based on current incident data. Further the travel times depend on the characteristics of the road networks in the regions and the environment conditions, e.g., weather and traffic congestion. None of these factors can be conveniently modeled by single variable functions of the region weights.

The workload metric has similar complexities. Since the actual workload depends on travel times, the complexities described above hold for workload. Additionally the types of incidents and their spatial distribution within and across the districts affect workload. The type of incident is also stochastic and cannot be calculated by single variable functions typical of districting problems.

The complex evaluation need for districting requires an approach that can produce accurate estimates of response time and workload for different districting designs. The next section describes our approach to the districting problem including a discussion of our evaluation methodology using simulation that is contained in Section "Evaluation of districting plans using simulation".

Approach to police patrol districting

The approach we have developed for the police districting problem begins by generating alternative districting plans. Subsection "Generation of alternative districting plans" gives details of the parameterized algorithm to generate the alternative plans. To evaluate these alternatives we have developed an agent-based simulation. We use the simulation to measure district performance using two metrics: average response time and workload. The simulation we use reveals important complexities in the choice of districting plans. More details on our simulation and its use for evaluating districting plans are contained in Subsection "Evaluation of districting plans using simulation". Because the simulation provides high fidelity evaluation of response times and workloads using data from actual CFS it is a time intensive process. So rather than evaluate all potential districting plans using the simulation we first evaluate some plans generated by randomizing the districting parameters. Then, the relationship between districting parameters and final performance measurements is studied and better districting plans can be generated. The iterative experiment-analysis-learning process can efficiently find better districting plans than random search method. We provide additional discussion of this screening step in Subsection "Relation between parameters and response variables".

Generation of alternative districting plans

The generation of districting plans is based on atomic geographical units. There are some existing geographical units such as police beats or census blocks. Usually, these geographical units consider administrative boundaries, important roads, or some natural boundaries (mountains, rivers). The redistricting procedure can start from these units and re-group them into several districts. When developing districting plans for large areas containing hundreds of such geographical units, police beats or census blocks are good choices for atomic units. However, some cities only have 20 or 30 census blocks. Police beats or census blocks are not small enough for optimal patrol boundaries. In such case, grid network can be

used instead. The city can be divided into several hundreds of grids and they are small enough to be atomic units. Clearly, more atomic units represent more possible districting plans. Such representation is more suitable for systematic and scientific study of districting problem. The output districting plans based on grid boundaries can be adjusted according to existing boundaries such as important roads, administrative boundaries and natural boundaries.

The main idea of our districting algorithm is similar to the Constraint-Based Polygonal Spatial Clustering (CPSC) algorithm developed by [13]. The two main steps of the algorithm are 1) select seed polygons (one atomic geographical unit) and 2) select the polygon in neighborhood to be added to the existing cluster until all units are assigned to districts. The major difference of our redistricting algorithm is the selection of seed for each district. Rather than totally randomizing the locations of the seeds, we locate the seeds on several concentric circles over the urban region (see Figure 1). The underlying intuition is that there is a general trend that downtown area of city has more population than suburbs. So the CFS incidents and crime events are also more likely to happen in central region of city. In the concentric circles model, central region of city has more seeds than suburbs and the patrol districts in downtown area tend to be small, which may facilities quick response and reduce workload variation. When selecting the seeds, the minimum distance between these seeds should be greater than a threshold to allow space for the districts to grow. Then, each seed alternates in acquiring adjacent atomic units until a stopping criterion is reached. The stopping criterion is that the sum of CFS probabilities of units in the district is greater than a user selected bound.

During the growing process, the districts or existing cluster of polygons select one adjacent unit to develop. If there are several alternatives, it chooses the one that can maximize the compactness score. Some randomness can be added to this procedure. The difference between random growth and compact growth can be seen in Figure 2. The district can randomly choose one adjacent unit to
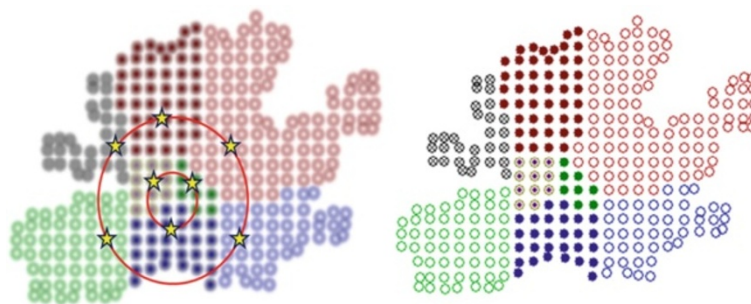

Figure 1 Seeds selection of police districting for Charlottesville.

Figure 2 Random growth vs. Compact growth.

develop. After a number of iterations of development, all districts stop growing because the stopping condition is met or because there are no more adjacent units for growth. Usually, there are still some unassigned units in the peripheral region of the city. They are assigned to the adjacent districts based on compactness principle. Then, the districting plan is adjusted to balance the CFS probabilities in each district and smooth the boundary between districts.

The locations of the seeds for districts determine the framework and basic structure of the districting plans. The relevant parameters for starting the growing process are: 1) the center of the concentric circles, 2) the number of circles, 3) the radius of each circle, 4) the number of seeds on each circle. Additional parameters determine the course of district growth. Examples are: 5) the stopping criterion, 6) the number of growth iterations, 7) growth randomness vs. compactness, 8) the number of iterations that balance the CFS probability between districts, 9) the number of iterations that smooth the boundary between districts. Therefore, a districting plan can be described and represented by a set of districting parameters. Once a districting plan is generated, some measurements can be quickly calculated without detailed simulation evaluation, such as compactness of plans and the variation of CFS probability of all districts. These intermediate measurements of districting plans can be used to select top proportion of plans for further simulation evaluation.

Relation between parameters and response variables

Without prior knowledge about how districting parameters affect final response variables, we randomize these parameters to generate some plans, quickly calculate the intermediate measurements, and take some time to get final responses through simulation evaluation. Then we build statistical models to study the relation between them, especially how districting parameters and intermediate measurements affect the final performance variables. The districting parameters can be adjusted to generate more plans that may have better performances. Since it is time intensive to use simulation to evaluate these plans,

they can be ranked by the combined weighted score of the intermediate metrics. The weights are adjusted based on the relation between intermediate variables and final performance variables.

Evaluation of districting plans using simulation

Because the assessment of response times and workloads requires the incorporation of multiple factors that interact in complex ways we cannot use closed form expressions. Also, field experiments in the law enforcement and safety management are clearly not feasible because of the risks and costs, not to mention, the public relations problems [14]. This means that evaluation of the police patrol districting plans requires a high fidelity simulation. A feature key needed in this simulation is the ability to accurately represent behaviors of the police in response to calls-for-service. Agent-based simulations afford the ability to effectively represent these behaviors.

Agent-based simulations capture of the behaviors of objects in an environment, such as police patrols in city, through the use of decision rules. These decision rules govern the interactions between objects in the simulation. For example, when a police car object interacts with a road object the rules specify the rate of transition to the next road object. These rules can also represent static properties of the object, for example the speed limit, and the dynamic properties of the environment, such as weather, construction, and traffic conditions. Other example rules used in our simulation include:

- The nearest police car to a new CFS responds to that CFS;
- The responding police car takes the shortest path to the location of the CFS;
- If a police car is servicing a CFS then it is no longer available to respond to a new CFS;
- If the nearest available police car is in a different district it will cross the district boundary and respond to the CFS.

The interactions between multiple objects governed by the rule sets in the simulation produce emergent

behaviors or properties that cannot be predicted before running the simulations. For our purposes the most important emergent properties are the response times to CFS and police workloads. These properties are the metrics that allow us to score the effectiveness of different patrol districting plans. Neither of these properties can be accurately anticipated a priori using only a districting plan and the numbers of CFS within the districts. As we indicated in Subsection "Evaluation complexity" if we could simply develop districting plans that equalize the expected CFS in each district then our problem reduces to the graph-partitioning problem. Although graph-partitioning is NP-hard; nonetheless, there are available heuristics than can be applied.

Instead the police districting problem has evaluation complexity (see Subsection "Evaluation complexity") in addition to the computational complexity of the graph-partitioning problem. In fact, when we use our high fidelity simulation to evaluate police districting plans that minimize the difference in CFS between districts they actually do worse than some other plans. The same is true for workload. The ability to discover these emerging properties is an important feature of the agent-based simulation we built and a critical requirement in the assessment of competing districting plans.

We implemented the simulation using Java Repast. Java Repast is an open source, agent-based modeling and simulation platform [15]. It uses object-oriented model and has a source library of classes for creating and running agent-based simulations and for displaying and collecting data from these simulations. Geographic data, such as the data expressed in shapefiles, can be imported into the Java Repast model. Using these geographic data the behaviors and movements of the agents can be controlled according to rule sets that exploit these data. For example, the shapefile data layer on secondary roads can have attributes that provide speed limits in different segments of this layer.

We built our police patrol simulation model as an extension to Malleson's RepastCity prototype [16]. The inputs to our simulation model consist of the shapefiles of the city, the patrol district plan, the police patrol allocation plan and a data set of CFS times, locations, and severity (this last attribute determines the distribution of the service time for the CFS). The shapefiles for the city include primary and secondary roads, major highways, and obstacles or impediments (e.g., construction or road work).

In order to simulate the time and spatial pattern of actual CFS and maintain the randomness of the city environment, the time between incidents and the locations are randomly chosen based on the distribution of actual CFS. Rather than use a bootstrap approach which would resample from actual CFS, we instead use an empirical fit of the distributions of the CFS in space and time. We then chose CFS values based on random draws according to these distributions.

Police cars are on patrol in their districts until they are dispatched by a CFS. Their patrol routes are randomly chosen from the network of roads in their assigned districts. An incoming CFS will generate the dispatch of the nearest police car and that car will follow the shortest route to the location of the CFS. In following this route it will use the maximum safe speed for the route which is greater than or equal to the speed limit on the route.

After the police car reaches the CFS location, the police car will remain at that location for the service time of that CFS. We obtain this service time as random draw from the service time distribution for the type and severity of the particular CFS. The service time distributions are empirical distributions found from the data set of actual CFS. When the service time ends the police car returns to its patrol route and again becomes available for dispatch to the next CFS.

We run the simulation for chosen number of runs or for a selected amount of simulation time. For each CFS in a run we record

- The time and location of the CFS;
- The identity of responding police car;
- The time of dispatch for the responding police car;
- The time of arrival at the CFS location for the responding police car;
- The time of departure from the CFS by the responding car (i.e., the time of arrival plus the service time);
- The travel distance of the responding car.

Using these data we can calculate the average response time and workload for each run and, hence, for each districting plan.

Case study and preliminary result

To illustrate the use of our approach to police patrol districting we used data from the Charlottesville, VA, USA police department. Charlottesville is a city with a diameter of about 7 miles and a population of about 40,000. However, this population increases during most of the year by another 26,000 due the presence of a major university. The population lives in multi-dwelling buildings, as well as, detached townhouses, apartments, and homes. There are more densely populated buildings near the university and the downtown. There are also commercial areas and some light industrial parks.

The current police patrol district of Charlottesville was designed more than 20 years ago. There are 8 patrol districts and in most of cases, one police car is assigned to patrol each district. The police managers and commanders want to draw district boundaries to incorporate

census block groups. These block groups are too large to serve as the atomic geographic units in our district growing algorithm (see Section "Approach to police patrol districting"). There are 37 block groups in Charlottesville. To create more useful atomic geographic units we decomposed the city into 323 grids. Figure 3 shows the locations of historical CFS incidents for several years, including 317,548 events. Many incidents happened at same places so each red point may represent many CFS events. To have a better view of the CFS density distribution in the city region, these historical incidents were spatially projected into the grid network. Based on the counts of CFS incidents, the CFS probability was calculated for each grid. The CFS distribution across these grids is shown in Figure 4.

As noted in Section "Approach to police patrol districting" our approach begins by generating thousands of plans using randomized districting parameters. Due to size of the city and number of districts, two circles are used to generate the seeds. For large cities, seeds can be located on

several circles, depending on the size of city and the number police cars. For cities with general shape of rectangular, ellipses can be used to generate seeds instead of circles. Then, these randomly generated plans were ranked by the weighted sum of two normalized intermediate measurements: standard deviation of CFS probabilities among districts and compactness score. Without prior information about the relevant importance of the two intermediate measurements of districting plans, they were considered equally important with the weights [0.5, 0.5]. Then, the top proportion of the ranked plans was evaluated using simulation model and final performance measurements were obtained. Due to the randomness of the districting algorithm, some combinations of parameter settings cannot generate compact districting plans. Furthermore, we cannot use simulation to evaluate too many districting plans due to the evaluation complexity of the problem, especially at the beginning phase of the experiment. In this case study, top 5% of the randomly generated districting plans were evaluated. More plans can be evaluated
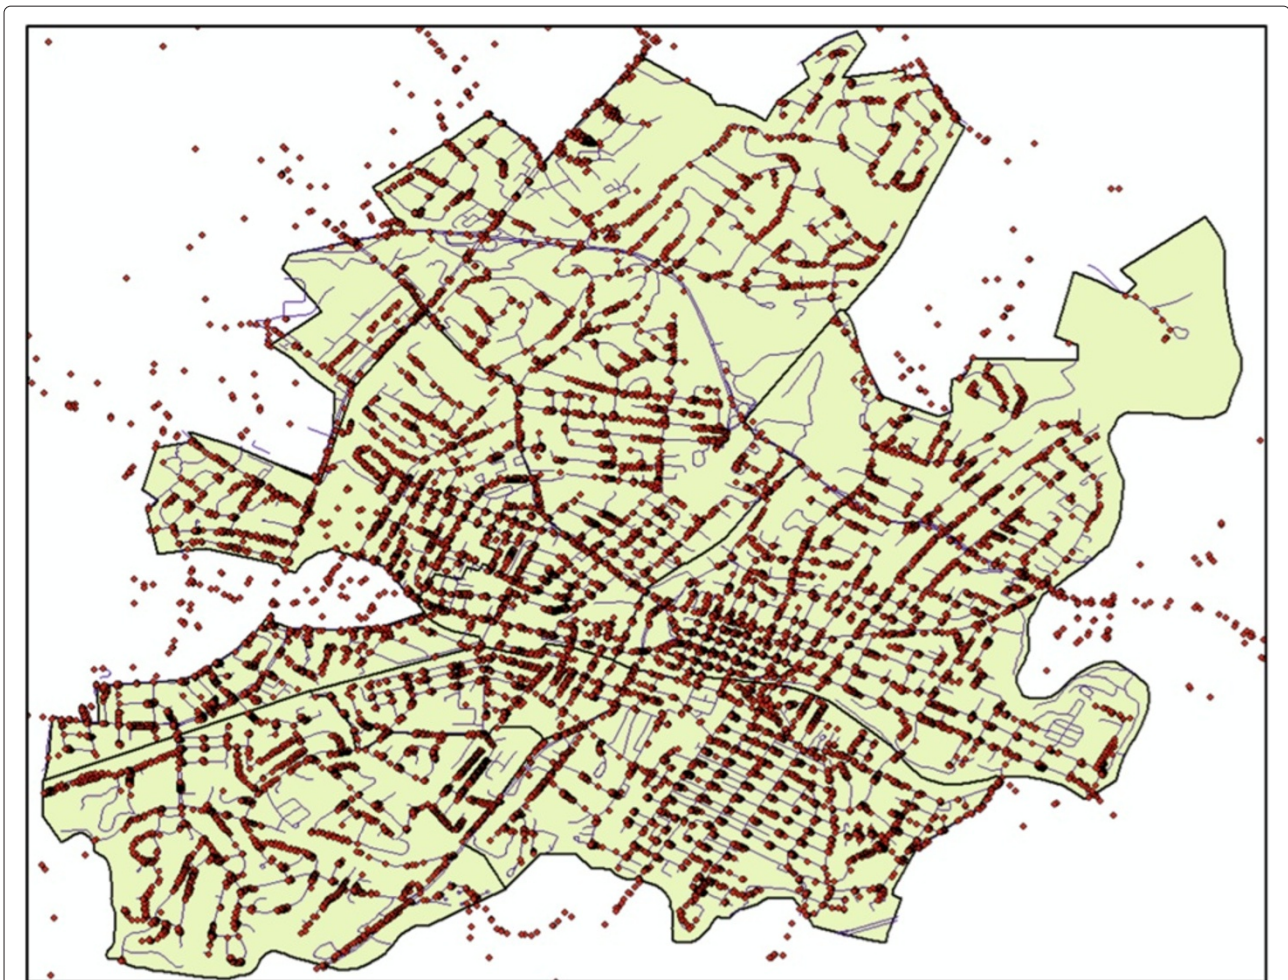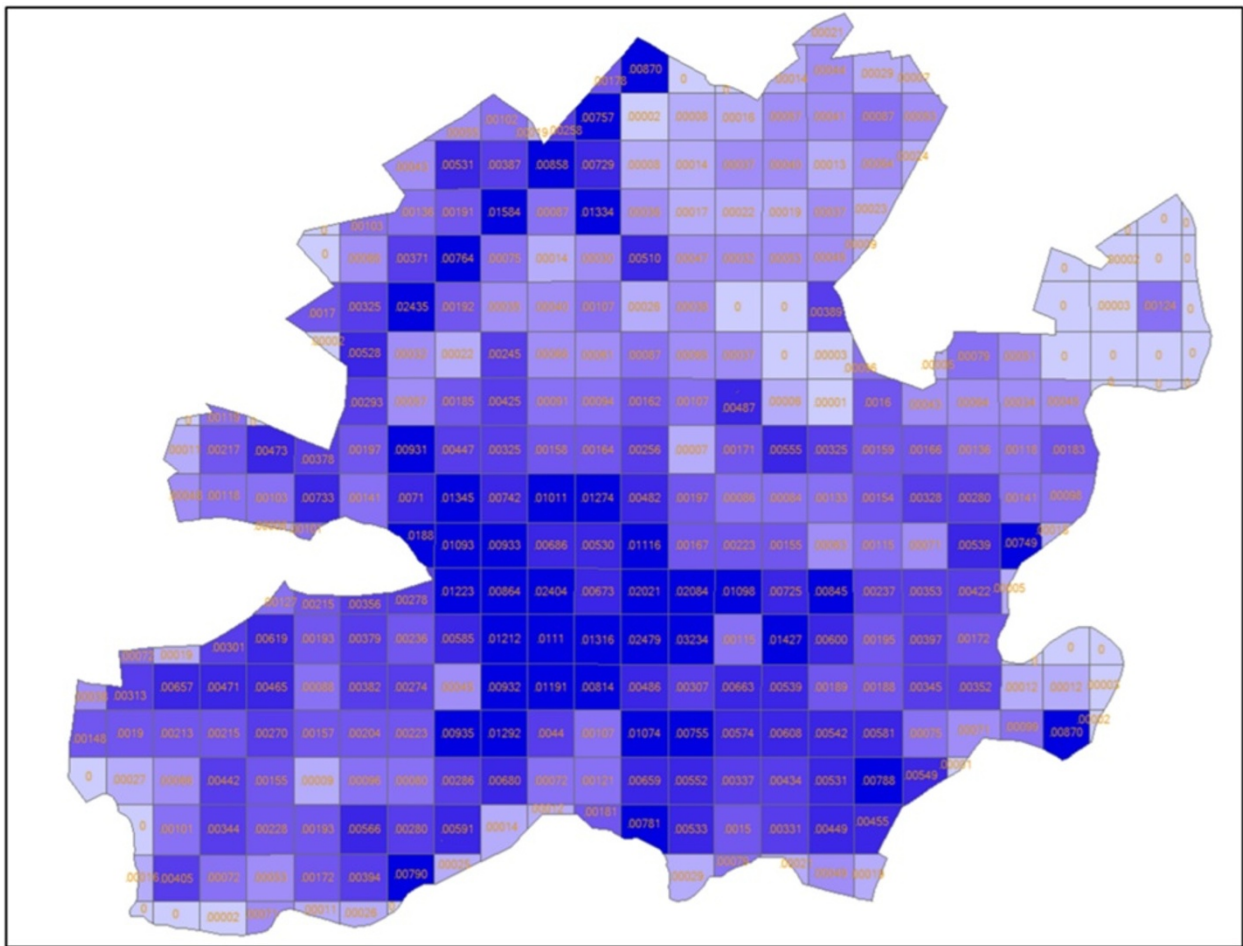

Figure 3 Locations of historical CFS incidents.

Figure 4 Call for service probabilities for 323 grids.
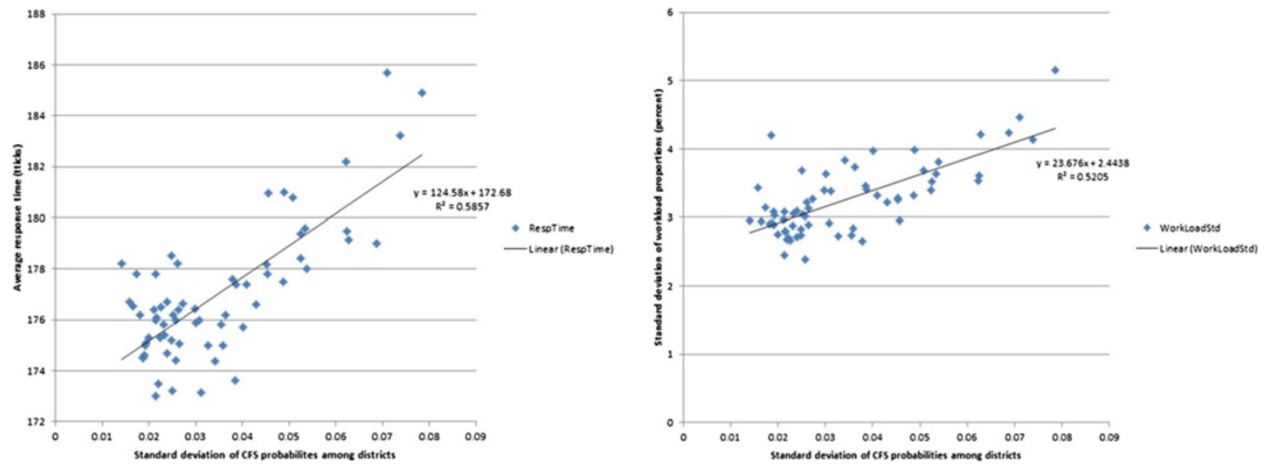

Figure 5 Linear relationship between Stdev. of CFS prob. and responses.

if more computational resources can be used. Based on result from statistical analysis, the center of the concentric circles, the radius of the outer circle are significant to average response time. For workload variation, the significant districting parameters are number of seeds on each circle, the center of the concentric circles and the number of iterations of balancing the CFS probabilities of districts. For the intermediate measurements, the variation of CFS probabilities is more important for both response variables. Its linear relationship to both responses can be seen in Figure 5. It can be seen that lower variation of cumulative CFS probabilities among districts leads to better performances for both responses. So, in the next iteration of ranking randomly generated plans, more weight can be given to standard deviation of CFS probabilities. Then, the relationship between responses and each districting parameter was analyzed individually and they were adjusted in the steepest decent direction to responses and another batch of thousands of plans was generated. The weights for two intermediate measurements became [0.4, 0.6] tentatively for compactness and standard deviation of CFS probabilities among districts. After ranking them

by the weighted sum, the top 5% plans were evaluated in-depth by simulation model. For each of the selected districting plans we ran the simulation for 450000 ticks (simulation time units) and 50 minutes of actual time. Under parallel running in current computing resources (13 batches on 3 PCs), the actual evaluation time for a districting plan can be reduced to 5 minutes. This length of time ensured the convergence of the simulation to a steady state, which can be seen in Figure 6.

Output analysis

While we put more effort into analysis of the performance metrics, average response time and workload, we also wanted to visualize the simulation. This allowed us and our police colleagues to visually assess and validate the behavior of the simulation. We showed the simulation to members of the police and they confirmed its behavior was consistent with that of their patrols. A static view of the simulation is shown in Figure 7. A dynamic view of the simulation can be seen in Additional file 1 or web link: http://www.youtube.com/watch?v=23ghFNAvdP4.
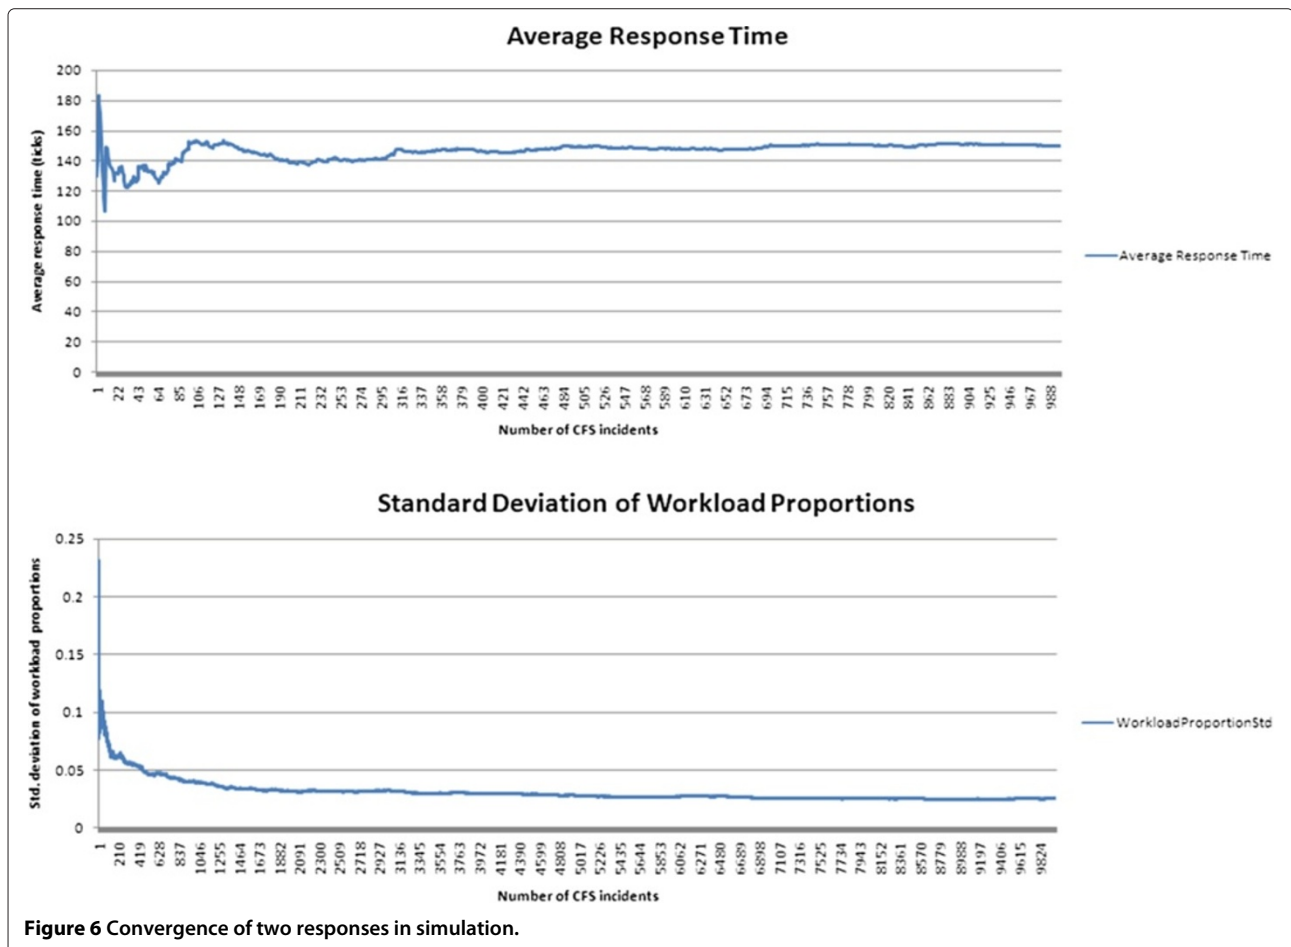

Figure 6 Convergence of two responses in simulation.

Figure 7 Visualization of Java Repast police patrol simulation model.

The current districting plan was also evaluated by the simulation system. The simulation results show both performance measurements can be improved. The average response time of current districting plan was 180.50 ticks (simulation time unit), which was reduced to 167.70 (Plan No.179). The standard deviation of workload proportion among 8 districts was reduced from 0.0340 to 0.0142 (Plan No.229).

Due to the NP-completeness of the graph-partition problem, there are too many possibilities of districting plans. We cannot use exact method to evaluate each of them. The evaluated districting plans in this case study are only a small proportion of the whole solution set and the solutions provided are preliminary. The global optimality cannot be guaranteed. We only find some significant districting parameters and intermediate measurements that may lead to better plans. More rigorous experimental design and statistical analysis can be conducted to further study the relationship between these factors and responses. With more powerful computational resources, more districting plans can be further generated and evaluated. It is possible to make improvements on both response variables.

Pareto frontier for bi-criteria decision making

Because we have a multiple objective problem (i.e., we want to minimize the average response time and minimize the variation in workload) we cannot choose a districting plan with a single objective. Further no single plan was best in both average response time and workload variation. To provide a multiple objective solution we used Pareto analysis. This analysis shows the positioning of each of the alternative districting plans with respect to each other on the two dimensional plot of both metrics. Using this plot we can trace the Pareto frontier which is the set of non-dominated districting plans. These plans are not dominated because no other plan is better than them in at least one of the performance metrics.
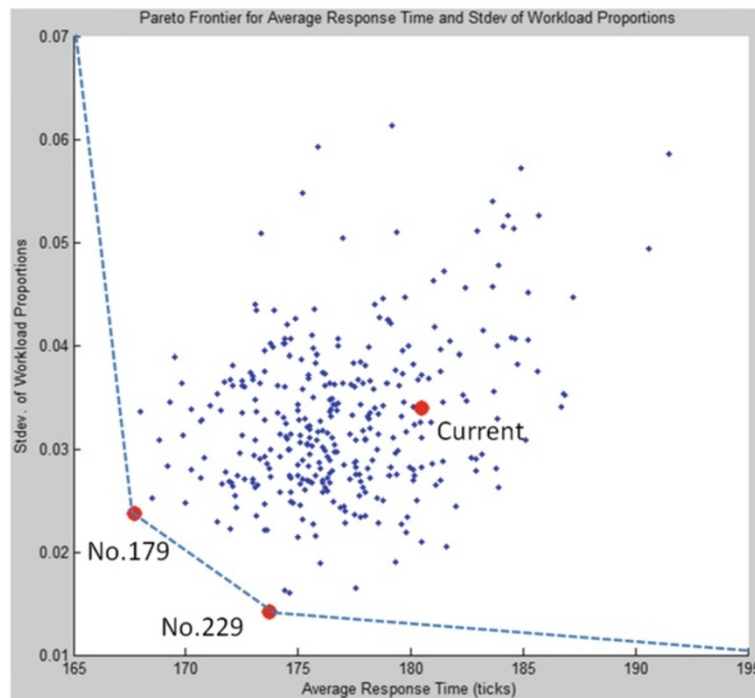
Figure 8 Pareto frontier for redistricting plans.

Figure 8 shows the Pareto frontier of average response time and workload standard deviation for the Charlottesville case study. This figure shows that 2 out of 300 districting plans are on the Pareto frontier. They are No.179 (best districting plan for response time), No.229 (best districting plan for workload variation). Under any weighting of response time and workload variation, one of the two districting plans dominates the others. So, they are the best 2 districting plans. The police department can choose one of them based on their needs and some practical considerations. The actual physical compositions of districting plan No.179 and No.229 are shown in Figure 9.

Replacing grid boundaries

As noted in Section "Generation of alternative districting plans", the police patrol boundaries should consider some administrative boundaries, natural boundaries and important roads. Clearly, the boundaries based on the grid
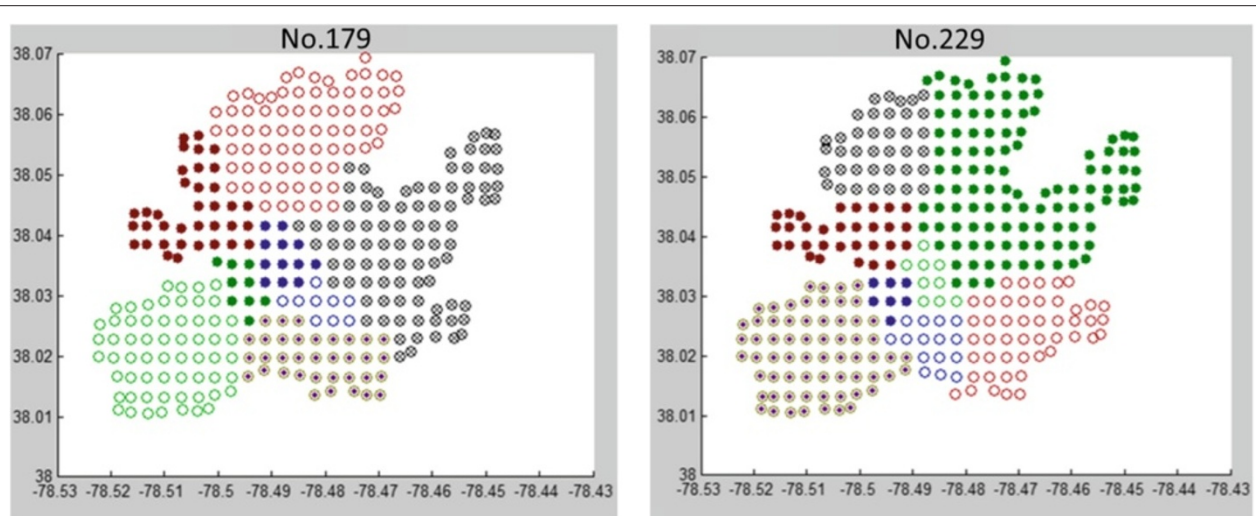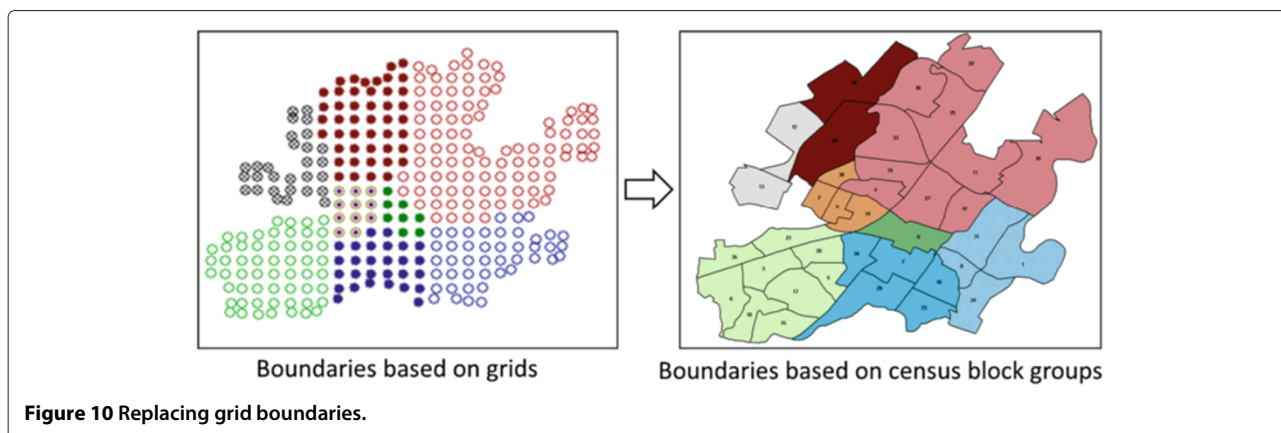

Figure 9 Districting plan No.179 and No.229.

Figure 10 Replacing grid boundaries.

network violate these boundaries. However, some existing geographical units such as police beats and census block groups consider these boundaries. So, replacing the grid boundaries needs to consulting with police departments. If patrol boundaries must be drawn based on police beats or census blocks, conversion can be made between grid network and these units. Example can be seen in Figure 10. If not necessary, the grid boundaries can be replaced by the nearest roads. In this way, the performance of the districting plan may be close the optimal solution based on grid network.

Conclusions and future work

In this paper, we reviewed the characteristics of the police patrol district design problem from the perspective of past and current work. We also showed the complexity of the problem from both computational and evaluation perspectives. In Section "Approach to police patrol districting" we described our approach to this problem. First, an algorithm for automatically generating police patrol districting plans was developed by growing them incrementally using small atomic geographic units as the building blocks. The second step is to screen these plans to obtain a smaller number of promising plans that can undergo more detailed and rigorous evaluation. Our last step is the perform evaluation using an agent-based simulation that provides high fidelity measures of performance. The case study presented in this paper shows the use of our police patrol districting approach for the Charlottesville Police Department. Compared with the current patrol district boundaries, the solution generated from our patrol districting approach shows improvement on both average response time and variation of the workload through the detailed simulation study.

While our simulation provides high fidelity it does not affectively capture changes over the course of an entire year. The changes include significant changes in weather and traffic patterns. To make our simulation in its current form work for an entire year would require run lengths

of several months. Hence we must develop improvements, either through sub-sampling, meta-modeling or other data reduction strategies. We also need to improve the screening method to ensure that we do not inadvertently remove an alternative that may prove superior in our more intensive simulation based evaluation.

Additional file

Additional file 1: Simulation dynamic view. This file is video record of Java Repast's visualization, which shows the dynamic view of the police patrol simulation system.

Authors' contributions
YZ implemented the simulation system, developed the districting algorithm in this paper, tested them, and drafted the manuscript. DB provided theoretical guidance in the whole procedure and revised the manuscript. Both authors read and approved the final manuscript.

References
1. M Rosenshine, Contributions to a theory of patrol scheduling. Oper. Res. Q (1970-1977). **21**, 99–106 (1970)
2. PE Taylor, SJ Huxley, A break from tradition for the San Francisco Police: patrol officer scheduling using an optimization-based decision support system. Interfaces. **19**, 4–24 (1989)
3. CD Hale, *Police Patrol, Operations and Management*. (Prentice Hall, Upper Saddle River, 1980)
4. SJ D'Amico, SJ Wang, R Batta, CM Rump, A simulated annealing approach to police district design. Comput. Oper. Res. **29**(6), 667–684 (2002)
5. PG Hancock, NC Simpson, Fifty years of operational research and emergency response. J. Oper. Res. Soc. **60**(51), 126–139 (2009)
6. GW Cordner, KE Scarborough, *Police Administration*. (Anderson, Cincinnati, 2010)
7. RC Larson, Measuring the Response Patterns of New York City Police Patrol Cars (1971). New York City Rand Institute R-673-NYC/HUD
8. PS Mitchell, Optimal selection of police patrol beats. J. Crim. Law. **63**(4), 577–584 (1972)
9. KM Curtin, K Hayslett-McCall, F Qiu, Determining optimal police patrol areas with maximal covering and backup covering location models. Netw. Spat. Econ. **10**, 125–145 (2010)

10. AAAly, DWLitwhiler, Police briefing stations: a location problem. AIIE Trans. (0569-5554). **11**, 12–22 (1979)
11. M Altman, The computational complexity of automated redistricting: Is automation the answer? Rutgers Comput. Technol. Law J. **23**, 81–142 (1997)
12. DS Johnson, The NP-completeness column: an ongoing guide. J. Algorithms. **3**(2), 182–195 (1982)
13. D Joshi, LK Soh, A Samal, in *Data Min, 2009. ICDM '09.Ninth IEEE. Int Conf*. Redistricting Using Heuristic-Based Polygonal Clustering, (2009), pp. 830–835
14. BA Knoppers, HF Miller, in *International Symposium on Criminal Justice Information and Statistics Systems Proceedings*. Computer simulation of police dispatching and patrol functions, (1972)
15. Official website of Java Repast. http://repast.sourceforge.net/
16. N Malleson, Agent-Based Modelling of Burglary, *PhD thesis*. (School of Geography, University of Leeds, 2010)