

SIMULATING CALLS FOR SERVICE FOR AN URBAN POLICE DEPARTMENT

J. Paul Brooks
David J. Edwards
Toni P. Sorrell
Sudharshana Srinivasan
Robyn L. Diehl

Virginia Commonwealth University
1015 Floyd Avenue, Room 4126
P.O. Box 843083
Richmond, VA 23284, USA

ABSTRACT

Police departments in the United States strive to schedule officers so that a number of benchmarks are met. The police administration is often asked to justify to local governing bodies the size of the police force. To assess the effects of force size and scheduling strategies on the ability to meet the benchmark goals, we develop a discrete-event simulation for the calls for service (CFS). Using actual call data from an urban police department in the United States, we fit distributions for call rates and service times for input to the simulation. The output of the model includes statistics related to the response delay, cross-sector calls, and officer utilization. The simulation model verifies intuitive notions about policing and reveals interesting properties in the system.

1 INTRODUCTION

A recent trend in policing urban areas in the United States has been to partition the precincts into smaller subregions called sectors. The local focus afforded by the sectors helps the departments meet and assess benchmarks for the policing system (Braga, Kennedy, Elin, and Peihl 2001, MacDonald 2002, Klein 2009, McGarrell, Corsaro, Hipple, and Bynum 2010). These goals include minimizing the response time to calls for service, providing for an equal workload for officers in different sectors, minimize the amount of overtime, and ensure that at least 40% of an officer's time is *proactive time*. Proactive time is the time spent patrolling the sector, engaging with the community, and servicing self-initiated events (e.g., when an officer observes a crime).

Operations research (OR) and emergency services such as policing have a long history, as described in the reviews by Green and Kolesar (Green and Kolesar 2004) and Simpson and Hancock (Simpson and Hancock 2009). As noted in both articles, publications in the area dropped sharply after 1989. Since that time, increases in computing power have enabled more in-depth analyses, but these are either not being reported in the literature or are not conducted.

In the early 1970s, several articles on policing, including analyses of travel time and allocation of officers to precincts, were published (Larson 1974, Chaiken and Dormont 1978, Larson 2002, Green and Kolesar 2004). In 1989, Taylor and Huxley (Taylor and Huxley 1989) won the Franz Edelman Prize recognizing excellence in the practice of OR for the development and implementation of a system for scheduling officers. Fazlollahi and Gordon (Fazlollahi and Gordon 1993) developed data-based approaches to solving crimes. More recently, Singer and Donoso (Singer and Donoso 2008) develop a queuing model for ambulance service which is applicable to police patrol cars. Curtin et al. (Curtin, Hayslett-McCall, and

Qiu 2010) describe methods for locating officers within precincts to facilitate faster response times. Much of the previous quantitative work analyzing policing is comprised of either deterministic math programming models for scheduling or queuing models that require unrealistic steady state assumptions. Our approach here is to depart from these frameworks and develop a simulation to characterize the system dynamics and learn more about how the policing system works under various conditions.

In this paper, we describe a discrete-event simulation for a police workforce assessment. We use actual data from the computer aided dispatch (CAD) system of a police department in a U.S. city to fit distributions for the interarrival times and service times of calls based on the time of day, the sector from which a call originates, and the priority level. The distributions are used as an input to the simulation, as well as a scheduling plan for officers. The simulation can be used to assess current practices and can forecast the effects of alternate scheduling strategies. The output of the simulation includes statistics concerning response times, cross-sector calls, and officer utilization.

The remainder of the paper is organized as follows. In Section 2, the raw data is described with the distributions that are fit and the benchmark goals for the police department. In Section 3, a discrete-event simulation is presented for analysis of the responses to calls for service. Section 4 contains results of the simulation analysis, including properties of the system that are revealed by the model. Section 5 synthesizes the results, highlights limitations of the simulation model, and describes future directions for this research.

2 DATA, DISTRIBUTIONS, AND BENCHMARKS

The data source for the workforce assessment simulation is a CAD system that is currently implemented to aid dispatchers in assigning calls to officers. The data can be interpreted in terms of the life cycle of a call for service (CFS):

1. Call is received by dispatcher.
2. Call is recorded.
3. Call is dispatched.
4. Officer arrives at scene.
5. Officer completes service, and is available for another call.

Each 24 hour day is split into six four-hour time periods. Calls are classified into one of three priority levels where the first level is assigned to the most urgent calls. Distributions of interarrival times and service times for calls are fit for each time period, priority, and sector of origin.

The interarrival time is the difference between the times that successive calls are received by the dispatcher. The service time is the time from when a call is dispatched to when the responding officer completes service, and is comprised of travel time plus *at scene time*. Distributions are fit for the *at scene time*, while estimates of travel time based on sector geography are used.

Distributions for the interarrival times and service times are selected based on Akaike's Information Criterion (see (Pawitan 2004)). Candidate distributions are exponential, Weibull, gamma, normal, and log-normal; the Weibull and log-normal distributions were dominant. Distributions are fit using the *fitdist()* function in the package *fitdistrplus* (Delignette-Muller, Pouillot, Denis, and Dutang 2010) for the R language for statistical computing (R Development Core Team).

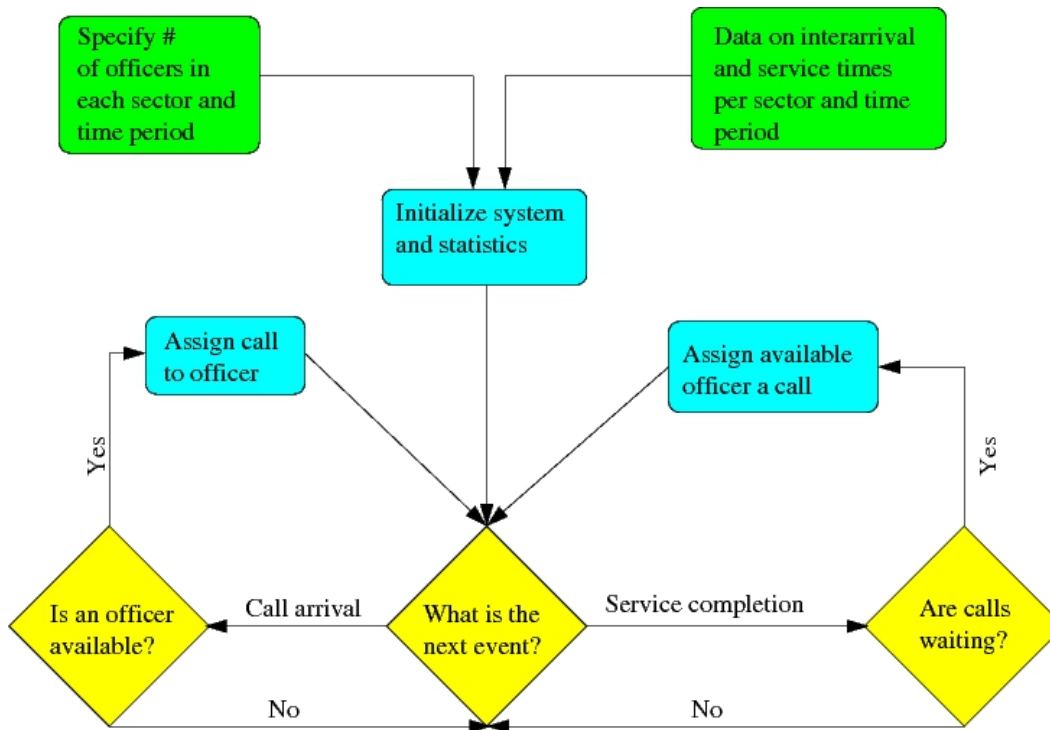


Figure 1: Flowchart of discrete-event simulation for calls for service.

Benchmark goals for the particular police department under consideration include:

1. Response times for priority 1 calls should be less than 7 minutes.
2. Calls should be serviced by officers from the originating sector; i.e., minimize cross-sector calls.
3. On average, officers should have 40% of their working hours as proactive time.
4. Minimize overtime hours.
5. The utilization of officers from different sectors should be approximately equal.

The response time is defined as the difference in time between when the dispatcher receives the call and when the first officer arrives at the scene. This time includes the time until an officer is available to take a call plus the travel time.

3 DISCRETE-EVENT SIMULATION OF CALLS FOR SERVICE

Figure 1 contains a flowchart that illustrates the discrete-event simulation for calls for service. The input information is comprised of distributions for interarrival times and *at scene times* for calls and a scheduling strategy for officers. The scheduling strategy specifies the number of officers working in each sector during each four hour time period of a day.

There are two types of events and a pseudoevent. Events are call arrivals and service completions, and the pseudoevent is the advancement of a time period. When the time period is advanced, the number of officers working is updated to reflect the schedule. Processing a call arrival involves assigning the call to an officer if one is available, or placing the call in the queue. If possible, a call is assigned to an officer in the originating sector. If not, then the call is assigned to an officer in another sector, with preference given to officers in sectors that are geographically nearby. The preference is determined by a predetermined ranking of the sectors for each originating sector based on proximity. If multiple officers are available in the selected sector, then the officer is chosen at random. Processing a service completion involves determining

if the officer is scheduled to continue working or if completing the current call required time past the scheduled shift. If the officer is still working and a call is in the queue, then a new call is assigned. Calls from that officer's home sector are assigned if possible; otherwise, the queues from other sectors based on the predetermined sector ranking are checked. The event calendar maintains the status of each officer (available/unavailable) and the time of the next service completion as well as the time of the next call arrival for each sector and priority.

The simulation is implemented using the C programming language. The random number generator of L'Ecuyer (L'Ecuyer 1999), implemented by Law (Law 2007) is used with a different stream assigned to each source of system randomness. For each configuration (officer scheduling strategy), the same streams are assigned to the same source of system randomness. Probability distributions are sampled by using a modified version of code distributed with the book by Leemis and Park (Leemis and Park 2006). The simulation is run for 30 simulated days for a warm-up period, and statistics are collected on a subsequent 7 simulated days. Statistics are collected over 35 replications.

4 RESULTS

To gain an understanding of how the policing system changes based on different scheduling strategies, the number of officers working in each sector during each time period is varied. For the results reported here, the number of officers working is the same across each time period and each sector. The number of officers per sector per time period is varied from 10 to 40.

The statistics collected include the average *response delay* for priority 1 calls for each sector and time period. The response delay is the time until an officer is available to respond to a call, and does not include travel time. The response delay serves as a lower bound on the response time, which is the subject of the first benchmark goal. The average number of cross-sector calls for each time period is also recorded for each combination of originating and serving sector. Finally, the percentage of four-hour periods for which *utilization* is less than 60% is calculated. Utilization is the proportion of time spent servicing calls, and so the percentage that we measure is a lower bound on the probability that the proactive time benchmark is met for a time period and sector.

Figure 2 contains a graph of the response delay for two sectors as a function of the number of officers per sector for each four-hour time period. The sectors, labeled Sector F and Sector K, are selected because they are the sectors with the most and least number of calls per service. In the figure, we see for the hours between 8am and 8pm, the average response delay is less than 10 seconds which agrees with the lower call rate during the day. For both sectors, the 12am to 4pm time interval has the highest response delays. The average response delay with 10 officers per sector is approximately 8 minutes for Sector F and 3 minutes for Sector K. As the number of officers per sector is increased, the average response delay drops rapidly, indicating that the only factor in determining response time under these conditions is travel time.

Figure 3 depicts perhaps the most interesting results from our simulation study. The figure contains three heatmaps for cross-sector calls for low (10), medium (20), and high (40) numbers of officers per sector. In Figure 3(a), most of the non-zero entries are along the diagonal, indicating that most calls are serviced by the originating sector. With only 10 officers per sector, most calls are queued, and when an officer becomes available, there is likely a call within his or her home sector waiting. In Figure 3(b), the system appears to reach a tipping point as cross-sector calls increase dramatically. With 20 officers per sector, there are just enough officers to cover the calls, as indicated by the low response delay times in Figure 2, but when an officer becomes available, there is a significant chance that all calls in the home sector are covered, and that the next call arrives in another sector. At this point, we see that the calls are likely properly covered, but system resources are poorly managed by sending units across the city to respond to calls. In Figure 3(c), all calls are answered by officers from the originating sector. With 40 officers per sector, notice that more calls are handled per time period on average than with 10 or 20 officers per sector. When an officer becomes available, there are usually no calls waiting, and when a call arrives, there is likely an officer available in the originating sector to respond.

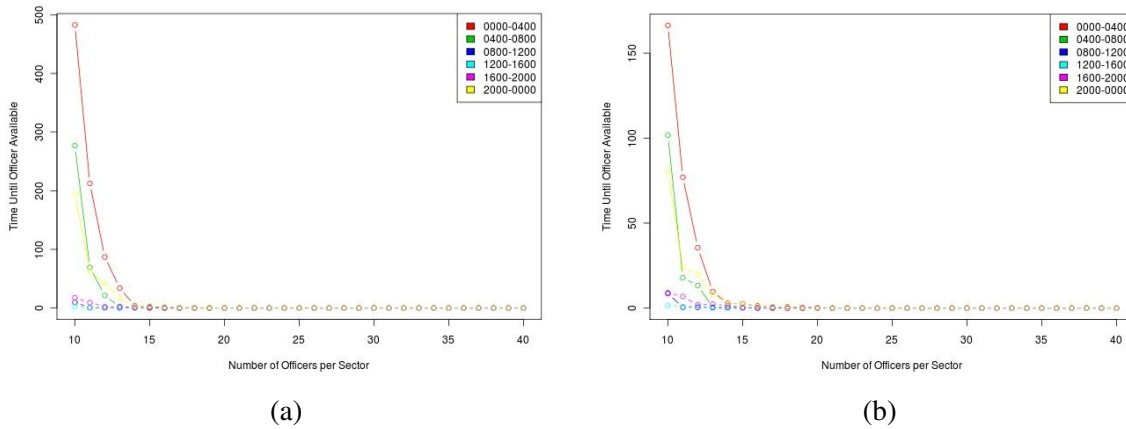


Figure 2: Response delay in seconds for priority one calls for (a) Sector F and (b) Sector K versus the number of officers per sector. A curve is plotted for each four-hour time period.

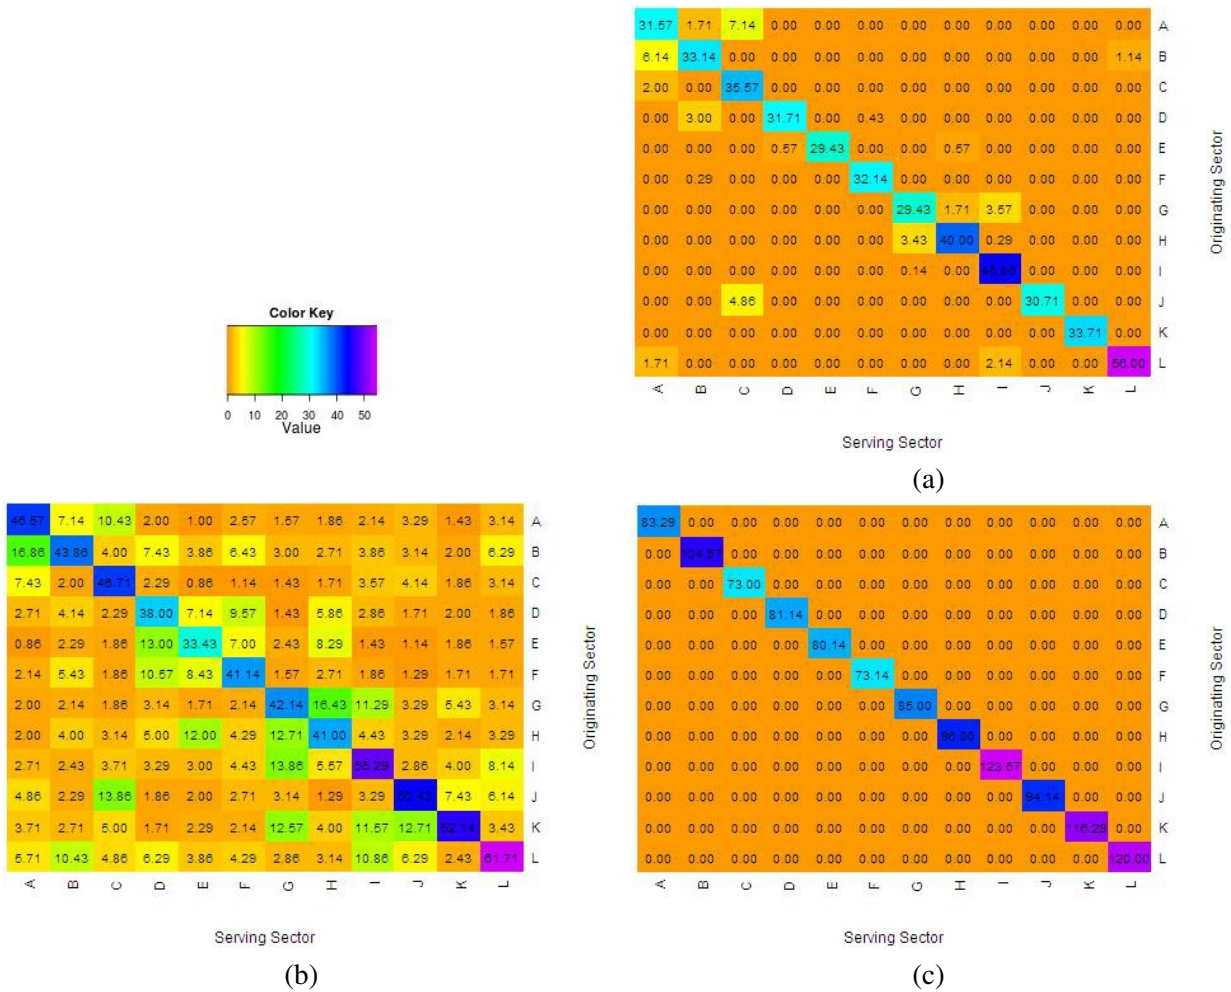


Figure 3: Heatmaps of cross-sector calls for (a) 10, (b) 20, and (c) 40 officers per sector for the 12am to 4pm time period. The number in each cell is the average number of calls originating in the call sector and answered by the serving sector.




Figure 4: The probability that officers spend less than 60% of each time period servicing calls for (a) Sector F and (b) Sector K versus the number of officers per sector. A curve is plotted for each four-hour time period.

Figure 4 contains plots of the probability that officers have less than 60% utilization in time period servicing calls for Sector F and Sector K as a function of the number of officers per sector. As expected based on the plots of response delay, the probability that officers have less than 60% utilization increases with officers per sector fastest for the time periods covering 8am to 8pm. Achieving the desired utilization in Sector F requires more officers than Sector K because of the increased volume of calls. For both sectors with 20 to 35 officers per sector, the probability of meeting the desired utilization is not monotone increasing, but is highly variable. This behavior is likely due to the effects of cross-sector calls which can affect utilization because of the additional travel time required.

5 CONCLUSIONS

We have presented results of a discrete-event simulation model of calls for service for an urban police department based on actual call data. The model is useful for revealing important properties about the policing system and it gives us insight into which benchmark goals are most difficult to meet. From the model described here, we can draw conclusions about the response time, the cross-sector call, and the proactive time benchmarks. It appears from our analysis that meeting the response time benchmark requires the fewest numbers of officers, followed by the cross-sector call benchmark and then the proactive time benchmark. If the number of officers per sector is just enough to meet call demand across the city, we observe a large increase in the number of cross-sector calls.

Several tasks remain before the workforce assessment can be completed. The model still needs to be properly validated against the raw data and summaries of the calls for service to ensure that anomalies in the data are not adversely affecting the distributions that are used in the simulation. In our model, we assume that one officer responds to each call, so that “officer” actually implies “unit”. The model needs to be updated to model the number of officers responding to each call. The geography represented in the model is at the sector level. The actual location of officers within a sector is not modeled, which hinders the ability to accurately capture the effects of travel time.

Future directions include using the simulation to forecast the number of officers needed in each sector during each four-hour time period in order to meet benchmark goals. A designed experiment and response surface methodology can help to determine the optimal configuration of officers. Once the best configuration is determined, an integer program can aid the scheduling of officers over two- and four-week horizons. Shift relief factors, based on actual data, will be used to estimate overtime hours and the size of the force needed to meet goals.

REFERENCES

- Braga, A., D. Kennedy, J. Elin, and A. Peihl. 2001. "Problem-oriented policing, deterrence, and youth violence: An evaluation of Boston's Operation Ceasefire". *Journal of Research in Crime and Delinquency* 38:195–225.
- Chaiken, J., and P. Dormont. 1978. "A patrol car allocation model: background". *Management Science* 24:1280–1290.
- Curtin, K., K. Hayslett-McCall, and F. Qiu. 2010. "Determining optimal police patrol areas with maximal covering and backup covering location models". *Networks and Spatial Economics* 10:125–145.
- Delignette-Muller, M., R. Pouillot, J.-B. Denis, and C. Dutang. 2010. "fitdistrplus: Help to fit of a parametric distribution to non-censored or censored data". <http://cran.r-project.org/web/packages/fitdistrplus/index.html>.
- Fazlollahi, B., and J. Gordon. 1993. "CATCH: Computer assisted tracking of criminal histories system". *Interfaces* 23:51–62.
- Green, L., and P. Kolesar. 2004. "Improving emergency responsiveness with management science". *Management Science* 50:1001–1014.
- Klein, G. 2009. "For the new Commander in Chief: A violence prevention strategy". *Journal of Police Crisis Negotiations* 9:55–60.
- Larson, R. 1974. "A hypercube queuing model for facility location and redistricting in urban emergency services". *Computers and Operations Research* 1:67–95.
- Larson, R. 2002. "Public sector operations research: A personal journey". *Operations Research* 50:135–145.
- Law, A. 2007. *Simulation Modeling and Analysis*. 4th ed. McGraw Hill.
- L'Ecuyer, P. 1999. "Good parameters and implementations for combined multiple recursive random number generators". *Operations Research* 47:159–164.
- Leemis, L., and S. Park. 2006. *Discrete-Event Simulation: A First Course*. Pearson.
- MacDonald, J. 2002. "The effectiveness of community policing in reducing urban violence". *Crime & Delinquency* 48:592–618.
- McGarrell, E., N. Corsaro, N. Hipple, and T. Bynum. 2010. "Project safe neighborhoods and violent crime trends in U.S. cities: Assessing violent crime impact". *Journal of Quantitative Criminology* 26:165–190.
- Pawitan, Y. 2004. *In All Likelihood: Statistical Modelling and Inference Using Likelihood*. Oxford.
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing, <http://www.R-project.org/>. ISBN 3-900051-07-0.
- Simpson, N., and P. Hancock. 2009. "Fifty years of operational research and emergency response". *Journal of the Operational Research Society* 60:S126–S139.
- Singer, M., and P. Donoso. 2008. "Assessing an ambulance service with queuing theory". *Computers & Operations Research* 35:2549–2560.
- Taylor, P., and S. Huxley. 1989. "A break from tradition for the San Francisco Police: Patrol officer scheduling using an optimization-based decision support system". *Interfaces* 19:4–24.

AUTHOR BIOGRAPHIES

J. PAUL BROOKS is an assistant professor in the Department of Statistical Sciences and Operations Research and Fellow of the Center for the Study of Biological Complexity at Virginia Commonwealth University. He received B.A. degrees in Mathematics and Physics from the University of Virginia, and an M.S. in Operations Research and a Ph.D. in Industrial Engineering from Georgia Tech. He is a member of INFORMS, the INFORMS Section on Data Mining, the INFORMS Computing Society, the Mathematical Optimization Society, and the Virginia Academy of Sciences. His research interests are in applications of operations research to biomedical research and to the public sector, and in the use of optimization for the development of data mining algorithms. His email address is jpbrooks@vcu.edu.

DAVID J. EDWARDS is an assistant professor in the Department of Statistical Sciences and Operations Research at Virginia Commonwealth University. He received a B.S. in Mathematics from Virginia Tech and an M.S. and Ph.D. in Statistics from the University of Tennessee. He is a member of the American Statistical Association, the American Society for Quality, and the Virginia Academy of Sciences. His research interests are in design and analysis of industrial experiments with especial interest in sequential experimentation and response surface methods. His e-mail is dedwards7@vcu.edu.

TONI P. SORRELL is currently an instructor for the Department of Mathematics at Virginia Commonwealth University and is pursuing her Ph.D. in Systems Modeling and Analysis. She received her Bachelor of Science in Secondary Education from Penn State University and was supported with an NSF Fellowship to complete her Master of Interdisciplinary Studies from Virginia Commonwealth University. She holds her National Board Certification in AYA Mathematics and is an IEEE Member. Her research interests include stochastic and Bayesian simulation. Her email address is tpsorrel@vcu.edu.

SUDHARSHANA SRINIVASAN is currently a doctoral student at Virginia Commonwealth University, working towards a degree in Systems Modeling and Analysis. She received her undergraduate degree in Computer Science Engineering from Anna University, India and her Master of Science in Mathematical Sciences with a concentration in Applied Mathematics from Virginia Commonwealth University. She is an INFORMS member, and her current research interests include scheduling, combinatorial optimization and stochastic simulation. Her email address is srinivasan3@vcu.edu.

ROBYN L. DIEHL is an Associate Professor of Criminal Justice in the L. Douglas Wilder School of Government and Public Affairs at Virginia Commonwealth University. She received her undergraduate degree from Randolph-Macon College in Psychology, her Master of Science from Virginia Commonwealth University in Criminal Justice, and her Ph.D. in Developmental Psychology from Virginia Commonwealth University. Her primary areas of research are program evaluation, project impact, and the impact of violent crime on youth and community development. Her research involves working with service-based organizations and law enforcement agencies to evaluate the efficacy of programs offered to residents in communities that suffer many of the negative consequences resulting from high levels of crime, specifically violent crime. Her email address is rldiehl@vcu.edu.