This article was downloaded by: [205.155.65.226] On: 10 July 2015, At: 13:48
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Casualty Collection Points Optimization: A Study for the District of Columbia

Aruna Apte, Curtis Heidtke, Javier Salmerón

To cite this article:
Aruna Apte, Curtis Heidtke, Javier Salmerón (2015) Casualty Collection Points Optimization: A Study for the District of Columbia. Interfaces 45(2):149-165. http://dx.doi.org/10.1287/inte.2014.0757

Full terms and conditions of use: http://pubsonline.informs.org/page/terms-and-conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2015, INFORMS

Please scroll down for article—it is on subsequent pages

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

Interfaces
Vol. 45, No. 2, March–April 2015, pp. 149–165
ISSN 0092-2102 (print) | ISSN 1526-551X (online)

http://dx.doi.org/10.1287/inte.2014.0757

Casualty Collection Points Optimization: A Study for the District of Columbia

Aruna Apte
Graduate School of Business and Public Policy, Naval Postgraduate School, Monterey, California 93943,
auapte@nps.edu

Curtis Heidtke
Office of Counsel, Naval Postgraduate School, Monterey, California 93943, cheidtke@nps.edu

Javier Salmerón
Operations Research Department, Naval Postgraduate School, Monterey, California 93943, jsalmero@nps.edu

A casualty collection point (CCP) is a single, predetermined location that is organized, staffed, and equipped to provide decontamination (if required), emergency medical assessment, treatment, and, where necessary, onward transportation of victims of a mass casualty incident. Emergency planners in the District of Columbia have recognized the desirability of developing a tool to assist planners in selecting CCPs within the affected area following a major incident. We develop a CCP optimization model (CCPOM) that provides planners and policymakers with strategic and operational insights into the complex problem of selecting optimal CCP locations to maximize casualty throughput for a range of incident parameters. Even more relevant, the CCPOM determines the utilization of personnel, decontamination units, and ambulances, providing planners with a general structure for resource allocation and signaling shortfalls that may lead to bottlenecks in casualty processing at the CCPs. District planners found many nonintuitive CCPOM results to be significant to their planning, programming, and budgeting efforts, and now consider the model's categorized resource utilization to be an integral part in updating District plans for both national special security event planning and everyday events.

Keywords: casualty collection point; mixed-integer optimization; disaster relief.
History: This paper was refereed. Published online in *Articles in Advance* September 5, 2014.

The District of Columbia (the District or DC) presents a unique operating environment for emergency planners. It is at once a city, state, and federal entity. By law, it is also the nation's capital and the seat of federal government. In addition, it is an economic, cultural, and transportation center of gravity in its own right. During the day, the District accommodates more than a million workers. In steady state, it is home to more than half a million permanent residents. It boasts many national monuments and icons, and its tourist population, which numbers tens of thousands day to day and 15 million annually, can swell to several hundred thousand during large-scale national and international events.

In practice, although the U.S. Congress retains the right to review and overturn city council actions, the mayor and city council exercise responsibility for the day-to-day management of city functions. Necessarily, this includes planning, preparing for, and responding to disasters that might pose a threat to the District. The agency charged with this mission is the DC Homeland Security and Emergency Management Agency. The primary planning document to implement this mission is the December 2008 district response plan, a sensitive but unclassified (SBU) source document not ordinarily available to the public.

Pursuant to this plan, the DC Department of Health (DOH) is tasked with establishing casualty collection points (CCPs) in coordination with the District's fire and emergency medical services department. Publicly available District planning documents, including the plan, do not clearly define the functions of a CCP, describe its organization, or set forth criteria for its location. In reviewing a draft report on survivor receiving centers and staging areas provided by the District, as well as the other planning documents provided by the contracting officer technical representative, we have arrived at the following working definition:

> *A casualty collection point is a single, predetermined location that is organized, staffed, and equipped to provide decontamination (as required), emergency medical assessment, treatment, and, where necessary, onward transportation of victims of a mass casualty incident.*

In the past, District planners selected CCPs using ad hoc unwritten criteria and nontechnical methodologies. However, the District has recognized the desirability of developing a strategic and operational tool to assist planners in selecting CCPs following a major incident. In this paper, we discuss our development of such a tool, which is a formal mathematical optimization model that seeks to maximize the throughput of victims needing initial healthcare assessments and, where necessary, treatment in the period immediately following a major incident. District planners validated the CCP optimization model (CCPOM) against a scenario developed by emergency planners based on a District threat assessment.

The result of this study is an operational tool that provides District planners and policymakers with strategic insights across a limitless range of incident parameters. Although we performed the current study for a single incident site (IS), we extended the model to investigate a broader utility, which lies in its ability to handle a larger number of hypothetical sites. The model also provides the planners with a framework for resource allocation and budgeting.

Related Literature

The use of optimization models to inform disaster planning and response is relatively recent, but not novel. Most of those studies focus on optimal facility locations. In our view, analysis of models such as *p*-median and set covering for location problems is quite mature in the literature. Brandeau and Chiu (1989) present a survey of a broad range of major location problems. Marianov and ReVelle (1995) review locating emergency services that involve spatial allocation. Owen and Daskin (1998) address the strategic facility location problems focusing on stochastic and dynamic characteristics that are integral components of emergency planning. Daskin and Dean (2004) classify location problems for health-services facilities in accessibility, adaptability, and availability models, and suggest applying concepts such as scenario planning to health facilities. The review by Caunhye et al. (2012) and the references therein provide further studies.

In a disaster relief environment and broadly in the public sector, the literature focuses on social welfare through equity-based objectives by minimizing the variability of the distribution of distances (Eiselt and Laporte 1995, Erkut 1993). Dekle et al. (2005) model an analogous location problem for Florida county disaster recovery centers as a set-covering problem, and use a two-stage approach to minimize the total number of centers. The maximal-covering location model in Balcik and Beamon (2008) incorporates inventory decisions in response to a sudden-onset disaster by determining the number and location of distribution centers and the quantity of inventory at each center. Lee et al. (2009) use a capacitated facility location model to determine points of dispensing for the Atlanta metropolitan area. The model considers population densities, maximum travel distance, available private and public facilities, and staff.

Drezner (2004) introduces CCPs, the primary motivation for our study. Drezner et al. (2005) carry this notion further by formally formulating a multiobjective location problem, which they illustrate with a hypothetical large earthquake striking Orange County, California. The authors propose a minimax regret multiobjective model. The individual objectives are based on the following concepts: *p*-median (average distance, e.g., Hakimi 1964, Daskin 1995); *p*-centre (maximum distance, e.g., Hakimi 1965, Daskin 1995); *p*-MaxCover (population covered within several radii, e.g., Church and ReVelle 1974); and minimum variance (distance equity, e.g., Maimon 1986).

Most of the earlier research develops location models with some type of equity measure as a primary goal; for example, they minimize the maximum distance between any individual and the facilities for distribution of emergency supplies, such as in the problem of dispensing vaccines that Lee et al. (2009) discuss. The existing research differs from our study in both scope and objectives: We seek to locate and resource triage facilities to optimize casualty throughput. In so doing, we also minimize the aggregate time necessary for onward transportation from CCPs to shelters. Our focus is on both locating the facilities and optimizing the allocation of resources required to establish those facilities. In this process, we identify the resources for which a potentially critical shortfall exists.

Scenario Overview

District representatives provided simulated conditions and supporting data to guide CCPOM development and desired outputs. Our study focuses on CCP placement recommendations based on a single test scenario: the detonation of a one-kiloton (1-kT) improvised nuclear device (IND) with Union Station as its epicenter (i.e., as the IS for the scenario). Although the given 1-kT IND scenario is not one of the 15 national planning scenarios (Federal Emergency Management Agency 2009), it is sufficiently similar to Scenario 1 (i.e., nuclear detonation—improvised nuclear device) of these 15 scenarios to warrant using the model for IND preparedness efforts.

The data provided can be classified into the detonation and its effect. The detonation data include: the nature of the detonation (ground burst); the relevant period of interest (48 hours following the incident); day of the week and time of the incident (weekday, daytime); and weather conditions. With assumed west-southwest prevailing winds, the fallout plume distributes east-northeast, away from the high-density areas west and south of the IS. We note, however, that a shift in actual winds could dramatically alter the number and distribution of victims and the availability of candidate CCPs, and hence the ultimate model results. The data on the effect of the detonation include: the number of casualties and the healthcare requirements of the victims (e.g., victims requiring triage and follow-on care versus the walking well); the speed distribution of casualties and the time when those casualties would start seeking medical treatment; and minimum and maximum radii (0.5 and 2.0 miles, respectively) for CCP location, expressed as a distance from the IS.

An IND incident of the type studied here could be expected to result in a nontrivial number of fatalities and critically wounded victims requiring immediate care. CCPs would not ordinarily process either category of victim because they will not perform mortuary functions and the critically wounded will be transported directly from the IS to hospitals. The remaining victims, generally termed casualties, are the sole subjects we consider in this study.

Structural damage resulting from a 1-kT nuclear detonation would be relatively minor. In terms of casualties, first responders and CCP staff can expect that the most significant healthcare issues will arise from victim radiation exposure after the initial blast. For our test scenario, we assume that 13,000 casualties will require medical care at CCPs, as predicted by a Defense Threat Reduction Agency analysis. We further divide these casualties into two distinct types (or waves) of victims arriving at the CCPs. The first type (i.e., first wave) includes those affected by the immediate blast. The second type (i.e., second wave) includes those who are affected by radiation from the nuclear weapon, but who do not realize they need treatment until after they experience symptoms.

In addition to scenario-specific data, the District has also provided data and criteria that are common to multiple disasters. Those include: descriptions and locations of available hospitals and shelters; healthcare worker data for CCPs (e.g., the average time required to triage a patient and the personnel required); availability of ambulances (including ambulance buses) and other forms of transport; average walking speeds for ambulatory victims; and a generalized list of evaluation criteria for determining the suitability of a particular location as a CCP.

CCP Prepositioning Optimization Model

Model Description

The CCPOM is a mathematical optimization model that primarily seeks to maximize casualty throughput at CCPs for a given set of conditions after an incident.

The CCPOM recommends the optimal location of CCPs, CCP sizing, staff (and other resource) levels, and transportation means allocation to maximize overall casualty throughput in the system. It determines which fraction of casualties from both waves should be directed to each selected CCP, and the casualty flow from CCPs to shelters and hospitals. Within the CCPs, the model accounts for casualty queues as patients waiting to be decontaminated, triaged, or otherwise treated.

CCPs are triage sites. They are neither intended nor equipped to provide long-term or specialized care to victims. Accordingly, the CCPOM assumes that critically ill or wounded casualties will be transported to hospitals after their initial evaluation and stabilization at a CCP. District data estimate these as approximately 10 percent of the casualties. CCPOM ambulance allocation to CCPs ensures adequate capacity to transport such casualties to hospitals.

Those casualties not requiring follow-on care, but who cannot return home (e.g., because of structural damage to their homes or radiation hazards), are assumed to walk or use public transportation to predetermined shelter locations. In addition to maximizing casualty throughput, as a secondary goal, the CCPOM also minimizes travel time from CCPs to shelters. Victims who reside outside the District but need temporary shelter (i.e., evacuees) are accommodated until they can arrange transportation to their homes. The remaining casualties require long-term shelter and subsistence. These displaced persons (e.g., those who live within the blast radius or, more likely, whose homes are in the projected plume path) will not be allowed to return home until conditions permit. Thus, casualties arrive at CCPs in a first or second wave, and leave CCPs as discharged, hospitalized, or sheltered as either evacuees or displaced persons.

Formulation

We worked closely with District representatives to refine and validate the data and specification of constraints that the optimization mode must enforce. We agreed that the solution produced by the CCPOM must ensure: (1) conservation of flow from both casualty waves at CCPs, shelters, and hospitals; (2) casualty flow management within CCPs based on the CCP's capacity for the various treatments; (3) staff and other resources allocated to CCPs cannot exceed prespecified levels; (4) shelter capacities for short- and long-term casualties cannot be exceeded; (5) allocated ambulance seat capacity must be adequate for transportation of some casualties to hospitals, accounting for travel times; and (6) hospital bed capacity cannot be exceeded. The appendix shows the mathematical formulation of the problem.

We implemented the CCPOM formulation in the General Algebraic Modeling System (GAMS Development Corporation 2013). The CCPOM applied in the baseline test scenario, which we describe next, has approximately 250,000 variables (104 binary) and 120,000 constraints, and runs in less than five minutes on a 2.4 GHz laptop to converge within a one percent gap.

Baseline Scenario: Data and Methodology

Planning Horizon and Other General Data

As established by the District, the planning horizon for the study begins at the time of the incident, referred to herein as T_0. The District's goal is to process all casualties as early as possible. District planners suggest 48 hours as an upper bound on the length of the planning horizon to accomplish this goal. Time is subdivided into intervals of 15 minutes (0.25 hours). This results in 192 distinct periods, where $t = 1$ corresponds with the first time interval after the incident (i.e., from T_0 to $T_0 + 0.25$ hours). Similarly, the last time interval ($t = 192$) corresponds to the last 15 minutes of our planning horizon (i.e., from $T_0 + 47.75$ to $T_0 + 48$).

The study posits a total affected population of 13,000 casualties. Of these, we assume 3,000 to be victims of the blast. These victims, the so-called walking wounded, will arrive in the first wave. The remaining 10,000 casualties, those suffering radiation effects, will arrive hours or days later in the second wave. The flow of victims through the system is modeled from the IS to CCPs, and then from CCPs onward to shelters or hospitals, or released outright. We derived the specific proportions and directions after consultations with District planners. However, to simplify the analysis, the model does not take into account increased fatality probabilities resulting from plume effects.

We expect that 40 percent of the casualties will be released to return home after being processed through a CCP. Of the remainder, 10 percent are expected to require medical evacuation to a local hospital, and 50 percent will need at least temporary respite at a designated area shelter. We further subdivide the number needing transport to shelters into 75 percent evacuees and 25 percent displaced.

Selection of Candidate CCPs

The District chose the blast site, IS, as ground zero for the IND detonation, and further specified that candidate CCPs must be no less than approximately 0.5 statute miles and no more than approximately 2.0 statute miles from the epicenter of the detonation. The rationale for the smaller of these parameters is that operating a CCP too close to the blast site would be too dangerous for responders for two reasons: residual blast effects and expectation of intolerably high radiation levels at that distance. The larger parameter reflects the planners' determination that expecting wounded victims to walk too far to reach a CCP is unreasonable.

Candidate CCPs are determined by using 100-grid squares of approximately 700 yards on a side. The grid is superimposed on an overhead satellite view of the subject area obtained from Google Earth (Google 2012). For all grid squares that are not *prima facie* infeasible, we use seven distinct criteria based on interviews with District emergency planners:

(1) Proximity: Reflects the distance of the candidate CCP from the IS, shelters, and hospitals.

(2) Access: Accounts for the location, size, and design of nearby roads, alleyways, driveways, and parking lots; the presence or absence of man-made barriers, such as fences, walls, lane dividers, and curbs; the presence or absence of natural barriers, such as hedges, trees, and berms; and the ability of the candidate site to accommodate large-scale decontamination equipment.

(3) Utilities: A fully functional CCP, particularly one with a decontamination capability, requires electrical power and water. Preferably, these would come from buildings or other structures located at or very near the chosen CCP site, which should have access to a fresh water source, either through a plumbing system or via fire hydrants.

(4) Space: Responders need adequate space to set up the various stations a fully operational CCP requires. A certain amount of physical space is also necessary for decontamination operations, site security, vehicle ingress and egress, and victim queuing.

(5) Fire stations: As the principal provider of emergency medical services (EMS) within the District, its fire department will have a significant share of the burden of choosing, setting up, and supporting the CCPs. Fire department equipment and personnel will be needed for mass decontamination, victim evacuation to hospitals, and other logistical support. CCPs will need to be rapidly established and supplied almost continuously. Locating the CCPs close to fire stations will facilitate this process. Of note, in its current form, the CCPOM does not factor in the distance from a candidate CCP to the nearest fire station in selecting CCPs, although we could easily add this to our model.

(6) Facilities: Physical facilities at a CCP site, or the absence of such facilities, can significantly affect its ability to accept large numbers of incident victims and CCP personnel. Certain broad generalizations can be made. For example, candidate sites with buildings are generally preferable to undeveloped sites during winter and inclement weather (e.g., a hurricane). The same may be said of sites with preexisting emergency equipment, such as electric generators or air and water filtration systems.

(7) Other considerations: Tactical, operational, and strategic factors that cannot be readily modeled may render a site unsuitable. For example, clustering two or more CCP locations in close proximity to each other may increase the danger of a secondary attack, or conflict with preplanned uses at specific CCPs may result in a situation that warrants restricting activity at a candidate site—or eliminating it as a viable alternative.

We used these criteria to assess each grid square for possible candidate CCPs. Where more than one suitable candidate CCP exists in a single grid square, we chose only one.

In addition to basic geospatial data from Google Earth, we used overlay files provided by the geographic information system-information technology branch of the District office of planning. These files depict location information for hospitals, schools, fire stations, community centers, and similar data. The previous process yielded 52 candidate CCP locations.

For purposes of calculating the CCP distance from ground zero, we used rectangular distance (instead of straight-line distance) as realistic in urban areas where people travel along established street patterns. The candidate CCPs ranged from 0.44 to 2.75 miles in rectangular distance from ground zero.

Shelters

The District Department of Transportation emergency transportation annex to the district response plan, also a SBU document, provides for the evacuation of victims of a mass casualty incident to a number of shelters dispersed throughout the District. The plan anticipates that these shelters will be used primarily as (1) way stations for evacuee victims awaiting transportation to destinations outside the District; and (2) longer-term shelters for displaced victims unable to return to their primary residences for extended periods.

In constructing the model, we assume that victims travel from any CCP to any given shelter on foot or by public transportation. The CCPOM determines the shelter to which it should send a victim based on the rectangular distance between the originating CCP and the shelter, the shelter's physical capacity, and the casualties being sent from other CCPs.

The emergency transportation annex of the District response plan identifies 21 specific shelter locations. After consultations with District planners, we added eight shelters for a total of 29 shelter locations.

In follow-up discussions, District planners proposed the potential inclusion of two other shelters of very large capacity, the National Guard Armory and the Convention Center. When we added these sites to the list of shelters, the CCPOM results were drastically skewed because it allocated all casualties to one of these shelters. We felt that these results were impractical. District planners also brought to our attention that the availability of either site to serve as a mass shelter is uncertain. Existing National Guard plans called for the Armory to be used primarily, if not exclusively, to support National Guard operations in the National Capital Region. However, the Armory's potential unavailability as a long-term shelter may render it suitable as a CCP, especially if Guard medical units are called in to assist civilian providers, or if the open space near the Armory is used for CCP operations. Therefore, we retained the Armory on the list of potential CCPs, but eliminated it as a shelter. In addition, the Convention Center is in a central location, close to many potential targets. During our conversations with District representatives, we discussed the issue of predisposing the model to choose one or both of these two sites over other suitable shelters. This appeared to risk trivializing operational considerations, such as multiple simultaneous or secondary attacks, including attacks on the shelters themselves, fallout patterns, competing uses, and the inherent logistical difficulties of sheltering many more victims in a single location. Accordingly, we excluded both locations from the list of shelters.

Hospitals

We reviewed hospital capability data for healthcare facilities in DC, Maryland, and Virginia. Because mutual aid arrangements with Maryland hospitals are in their early planning stages, we excluded Maryland facilities from consideration.

We derived the baseline information requirements from District of Columbia Hospital Association (2006), an internal study. The study evaluates the state of hospital readiness within the National Capital Region on a number of criteria that assess both infrastructure and capabilities. The study gathers information on hospital bed space, including surge capacity. Of the 27 hospitals included, 18 are surveyed. For the remaining nine hospitals, DC health emergency preparedness response administration (HEPRA) planners provided data. Five hospitals do not report surge capacity. Therefore, the CCPOM considers 22 hospitals that would meet the minimum acceptable standard of care for providing clinical services to arriving casualties.

Using information provided by the District, the CCPOM assumes that ambulances would travel from CCPs to hospitals and back at an average speed of five miles per hour for hospitals within the District. This low speed reflects expected delays for large-scale vehicular-traffic congestion and pedestrian flows away from the affected area. For hospitals in Virginia, it assumes an average speed of 15 miles per hour, reflecting somewhat lessened vehicular-traffic congestion and significantly reduced pedestrian traffic outside the District. Based on historical data, HEPRA planners determined that approximately 178 spaces (seats and litters) would be available, on average, for transporting victims between CCPs and hospitals. The CCPOM also assumes that each ambulance will be assigned to a unique CCP, but could transport casualties to any hospital, as needed.

Casualty Arrival at CCPs

District planners assume that the vast majority of victims would arrive at a CCP in one of two principal waves. The first wave (blast casualties) is expected to occur shortly after the primary IND explosion. The second arrival wave (radiation casualties) is expected to begin several hours after the incident and continue for a period of days. Drawing from lessons learned from past incidents and exercises, the expectation is that certain numbers of asymptomatic casualties will report to CCPs merely because they are in the general vicinity of the IS at the time of the incident and fear they might have been affected.

District planners provided invaluable operational perspective and feedback in developing the arrival distributions. The model assumes that the first casualty wave will not begin the journey from the IS to a CCP until the CCPs are operating. Based on planner inputs, the model expects that the first CCP will not be fully operational until two hours after the incident has occurred.

Once notified that CCPs are operational and ready to accept casualties, at approximately $T_0 + 2$ hours, emergency personnel at the IS will begin grouping casualties and directing them to specific CCPs. All casualties will not travel at the same speed. The CCPOM assumes an underlying, deterministic arrival pattern that District planners suggested based on three categories of distances from the CCPs to the IS: near (less than one mile), medium (between one and two miles), and far (more than two miles). Within each category, these deterministic patterns are such that their average travel times are proportional to the distance. Also, the arrival distribution patterns are more dispersed for casualties who travel longer distances (see Figure 1). For example, for a CCP located at 0.78 miles from the IS, the assumed deterministic arrival distribution has an average $\bar{x} = 2.94$ hours. This figure reflects two hours that casualties are assumed to remain at the IS plus an average of 0.94 hours walking to the CCP. Standard deviation (σ) is 0.50 hours. For a CCP located at 2.2 miles, these values are $\bar{x} = 4.65$ hours and $\sigma = 1.57$ hours, respectively.

A degree of uncertainty is inherent in estimating the civilian population's response to an IND incident of the sort hypothesized. In the absence of empirical data from past incidents, we relied on our discussions with experienced emergency planners. From these discussions, we derived that the onset of symptoms of the affected population (i.e., citizens who did not realize they were exposed) in the second wave occurs several hours to days after radiation exposure. This is consistent with the observed phases of acute radiation syndrome, in particular the onset of prodromal symptoms (Federal Emergency Management Agency 2010, p. 83). Other than the decontamination units, the CCPs are presumed not to have resources for explicit radiation treatment. Necessarily, such victims are sent to hospitals. Because reliably estimating the percentages in the first and second waves is difficult, we followed the guidance of District planners for casualty proportions in each wave.

Figure 2 shows the second-wave arrival pattern. Because the second wave of casualties is comprised almost exclusively of those suffering from delayed symptoms of radiation exposure, these individuals do not travel directly from the IS to the CCPs. As a practical matter, determining the point of origin for each casualty in this second wave is effectively impossible. To simplify the analysis, we assume that, regardless of point of origin, District emergency responders will direct these second-wave victims to an appropriate CCP. This is possible, for example, if we assume that

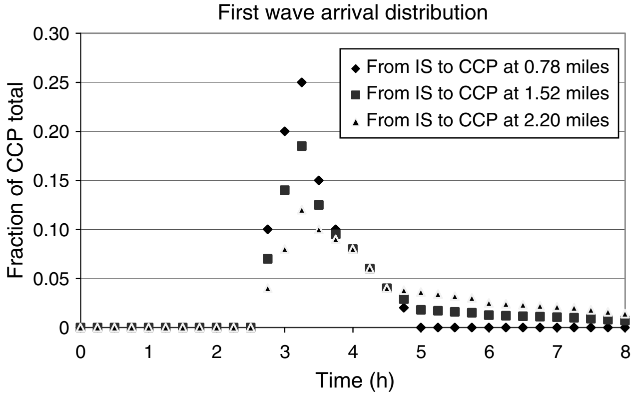

Figure 1: The diamonds, squares, and triangles show fractions of arrival times for first-wave casualties for sample CCPs at various distances from the incident site.

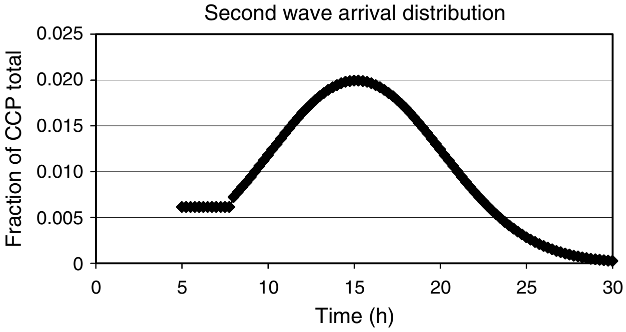

Figure 2: The arrival pattern of fractions of second-wave casualty arrivals at a CCP over time is close to normal distribution with the average arrival time occurring at approximately $T_0 + 15$ hours, as provided by District planners.

Resource (units)	Available	Minimum required (units/CCP)	Average time to process (min/casualty)	Casualties requiring the service (%)	Unitary capacity (casualties/unit-hour)	Total capacity (casualties/hour)
Decontamination units (units)	25	1	0.50	100	100.0	2,500
Triage EMS (persons)	72*	12	3.00	100	16.7	1,202.4
Administrators (persons)	300	3	2.00	100	25.0	7,500
EMS nurse supervisors (persons)	150	15	12.00	100	4.2	630
Behavioral staff (persons)	60	3	5.00	1	1, 000.0	60,000
Security, command, and control (persons)	112	12	2.00	100	25.0	2,800
Transportation preparation (persons)	30	3	2.00	1	2,500.0	75,000
Ambulance spaces (seats)	178	1	Depends on trips	10	Depends on trips	

Table 1: The data in the table are based on the input provided by the District planners.

***Includes resources projected to be supplied by National Medical Response Teams (NMRTs). The number of available medical staff includes five 12-person teams sourced from District healthcare facilities and 12 personnel from the NMRTs.**

these casualties call into a dispatching system to ask for guidance. Directly or indirectly, emergency responders can then retain significant control over the travel patterns of both waves of casualties.

Service Allocation to CCPs

Once they arrive at a CCP, casualties must wait to be processed, evaluated, treated, and released, or designated for follow-on transport, as appropriate. To operate efficiently and minimize the total time for patients to be treated, each CCP must be staffed adequately. Therefore, the CCPOM must also optimize the distribution of personnel, equipment, and other resources across all selected CCPs.

To accurately model internal CCP operations, we researched the notional makeup of a typical CCP and consulted with HEPRA planners. We deemed the following categories of resources to be required within each CCP: decontamination units, triage medical staff (primarily nurses and EMS personnel), administrators, EMS nurse supervisors, behavioral specialists, security personnel, command and control personnel, ambulance spaces, and transportation preparation teams.

To enhance the model's realism, we incorporated certain operational considerations into the design. First, the CCPOM broadly compensates for workload and personnel fatigue by assuming that all CCP personnel will work an average of 50 minutes per hour (i.e., 5/6 of the time). Given resource constraints, to assume that replacement personnel will be available for each CCP worker is unrealistic; therefore, the CCPOM does not make this assumption.

Second, the model assumes that most services are applied to 100 percent of the casualties; the exceptions are behavioral staff, which it applies to one percent of the casualties (usually because of mental health crises after the incident), and transportation preparation, which it applies to one percent of the casualties (a fraction of those casualties who go to hospitals also need special assistance).

Table 1 summarizes the essential characteristics of each resource. For example, the CCPOM assumes that triage EMS is 72. The minimum established medical staffing requirement at each CCP is 12 staff members. It takes three minutes, on average, to triage a casualty. Thus, hypothetically, each medical staff person allocated to a CCP could treat 20 casualties per hour. Because medical staff personnel are only available 5/6 of the time, we correct the hourly throughput for this service, which becomes approximately 16.7 casualties per triage EMS per hour, and list this number as unitary capacity. The last column reflects total hourly processing capacity for the resource.

Results

Baseline Scenario

For practical reasons, the District specified that the model should choose at least two and no more than five CCPs. When applied to the baseline scenario described in the previous section, the CCPOM selects five CCPs from the candidate list.

CCPOM choices do not show a clear pattern (see Figure 3) based on geography, size, or other obvious

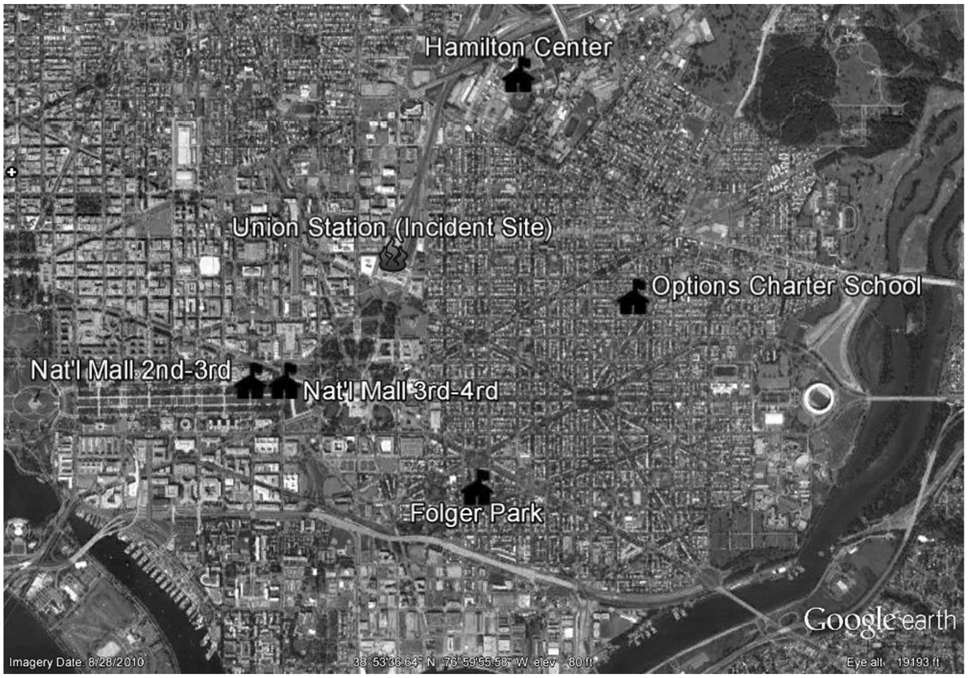

Figure 3: The house icons are the locations of optimal CCPs as generated by the CCPOM for the baseline scenario using Google Earth (Google 2012); the fire icon represents the incident site.

factors. The model clearly weighs a number of variables and complex relationships to arrive at a nontrivial solution.

Table 2 breaks down the number of casualties serviced by each CCP for each wave of arrivals. The proportion of second-wave victims to first-wave victims treated is roughly equivalent across CCPs. However, this does not imply that each CCP will service a proportionate share of casualties. Rather, as the results show, each CCP will service a different percentage of casualties based on factors such as distance from the IS, staffing levels, and ambulance availability.

CCP	First wave	Second wave	Total
Options Public Charter School (OPCS)	503	1,687	2,190
Folger Park (FP)	697	2,295	2,992
Hamilton Center (HC)	1,204	4,050	5,254
Nat'l Mall 3rd/4th st. (NM-3/4)	298	983	1,281
Nat'l Mall 4th/6th st. (NM-4/6)	298	985	1,283
Total	3,000	10,000	13,000

Table 2: For each wave of arrivals, the table lists CCPs selected in the baseline scenario and the distribution of casualties each CCP services.

The CCPOM projects all casualties can be treated and sent to their final destinations within the 48-hour planning horizon. Interestingly, the model projects that NM-3/4 and NM-4/6 will serve virtually identical numbers of casualties from both waves. Intuitively this makes sense, because both CCPs have identical characteristics and are adjacent to each other. Running independent operations with separate staffs, administrators, logistics, and command and control centers might lead to inefficiencies, competition for scarce resources, and congestion arising from unsynchronized pedestrian and vehicular-traffic flows. Conversely, creating a mega-CCP with a single command element may exceed the site commander's span of control, thus introducing bureaucratic confusion and inefficiencies. Ultimately, District decision makers must weigh the alternatives using input from analytical tools, such as the CCPOM, and other relevant sources, such as exercises, subject matter experts, and published best practices.

In addition to selecting the optimum CCP locations, the CCPOM allocates available staff and transportation resources among the selected CCPs. Table 3 shows the

	CCP					
Resource (units)	OPCS	FP	HC	NM-3/4	NM-4/6	Used resource (vs. available)
Decontamination units (units)	2	2	3	1	1	9 (25)
Triage EMS (persons)	12	12	15	12	12	63 (72)
Administrators (persons)	4.2	5.8	10	3	3	26 (300)
EMS nurse supervisors (persons)	25.4	34.6	60	15	15	150 (150)
Behavioral staff (persons)	3	3	3	3	3	15 (60)
Security, command and control (persons)	12	12	12	12	12	60 (112)
Transportation preparation (persons)	3	3	3	3	3	15 (30)
Ambulance spaces (seats)	25.7	46.8	74.0	14.7	16.8	178 (178)
Maximum throughput (casualties/hour)	105	144.3	250	62.6	62.6	

Table 3: The table lists the resource utilization in the baseline scenario for the CCPs. The last column compares the number of resources that are allocated to all CCPs with the amount of resources that were originally available for each resource category.

utilization of resources among the five CCPs. Given that HC has the greatest number of casualties processed (from Table 2), it requires more resources than the other CCPs: its maximum throughput of 250 casualties served per hour is also the largest among the chosen CCPs. That maximum throughput is capped by the individual resources available; it is calculated as the minimum of the product of each individual resource throughput and its associated throughput rate from Table 1.

We observe that EMS nurse supervisors are a critical resource, because all available personnel are utilized. The model allocates all available ambulance seats among the five CCPs, thus making this resource also critical. The combined effects of fully utilizing ambulance seats are to expedite the transportation of patients to hospitals with open bed space and minimize the time casualties need to wait for ambulances after being treated at a CCP. Note that less than 10 percent of the available administrative personnel is utilized. The number of available security and command and control personnel is almost double the number needed. We also observe that only nine of 25 decontamination units are utilized.

The model projects that nearly all shelters will be utilized to their maximum capacities for both displaced casualties and evacuees. The utilization results are based on one wave of arrivals at shelters. Operational experience may warrant post-study modifications to the model to account for additional waves, or to incorporate arrival patterns of a stochastic nature.

The final flow modeled by the CCPOM is the movement of more severely injured or ill victims from CCPs to hospitals. This accounts for 10 percent of all casualties treated at CCPs (1,300 victims).

Of the 22 hospitals included in the study, only INOVA Fairfax Hospital is not utilized, possibly because it is located farther from the IS. This reflects the attempt to reduce total transportation time of all ambulance trips, given the limitation of that resource (see Table 3). Of the hospitals receiving casualties, six are utilized to their full surge capacity. This occurs in spite of the hypothesis that a relatively modest number of casualties will need hospitalization. A larger-scale incident with a higher number of casualties or hospitalization rate could well tax regional medical assets beyond their capabilities.

The primary objective (z_1 in the appendix) measures promptness in casualty treatment by maximizing weighted throughput of casualties, where weights are inversely proportional to the time when casualties are treated (patients treated over time). Because each hour has four 15-minute periods, a casualty treated in period $t = 10$ (2.5 hours after the detonation) contributes $4/10$ ($= 0.4$) weighted casualties per hour to the objective function, whereas a casualty treated in period $t = 30$ contributes $4/30$ (≈ 0.13) weighted casualties per hour. These figures are then aggregated for all 13,000 casualties to yield a single number representative of overall casualties treated per hour. Overall, the objective function value for the baseline scenario is $z_1 = 1{,}390$ weighted casualties per hour. Note this does not imply an average of $13{,}000/1{,}390$ (≈ 9.35) hours per casualty

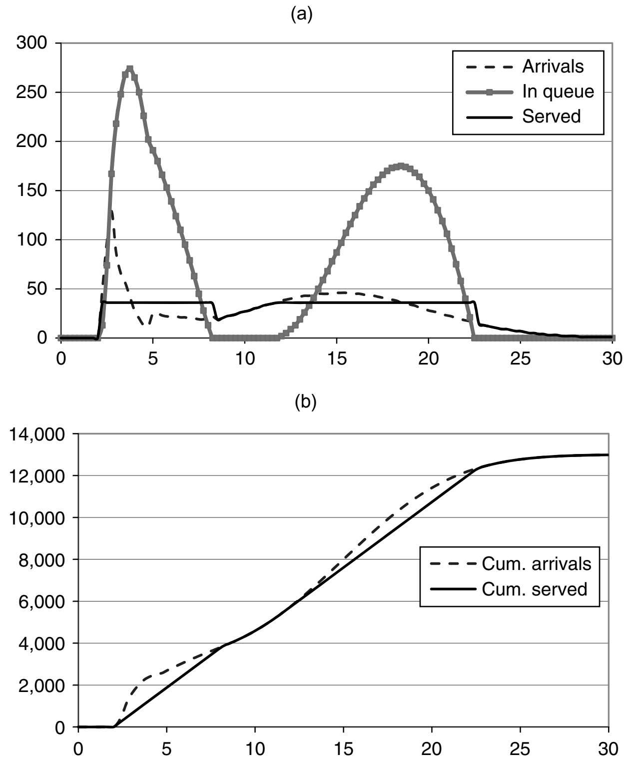

Figure 4: Figure (a) depicts the flow of patients from the CCP from Folger Park; (b) shows the cumulative flow for all CCPs in the baseline scenario. The horizontal axes show time (hours) and the vertical axes show number of casualties.

because of the weights used (inversely proportional to time). We can calculate the actual average time to process a casualty by using results on the number of casualties served by period and CCP (S_{ct} in the appendix). The resulting average time, including both waves, is 13.2 hours.

Figure 4(a) demonstrates the fluctuation of the queue length relative to arrival patterns and processed casualties at a selected CCP. The pattern evidenced in the figure is representative of those at other CCPs. The differential between total arrived versus served casualties is represented by the gap between the two graphs in Figure 4(b).

The second objective is the total time to shelters (z_2 in the appendix). This is simply the sum of all trip times incurred by 50 percent of the casualties as they travel from CCPs to shelters. Casualties in the baseline scenario require 8,242 hours.

CCP	First wave	Second wave	Total
NM-3/4	1,211	3,982	5,193
NM-4/6	1,212	3,974	5,186
Lafayette Square (LS)	577	2,044	2,621
Total	3,000	10,000	13,000

Table 4: For each wave of arrivals, the table lists the CCPs selected in the mall scenario and the distribution of casualties that each CCP services.

Mall Scenario

In this scenario, the CCPOM is constrained to consider only 11 candidate CCPs in the immediate vicinity of the National Mall. The CCPOM chooses three CCPs as optimal, including NM-3/4 and NM-4/6, which are more extensively used than in the baseline scenario. Table 4 summarizes the throughput results.

As in the baseline scenario, the CCPs chosen generate sufficient capacity to process all 13,000 casualties. Under this scenario, each NM CCP processes four times more casualties with respect to the baseline scenario. The CCPOM also treats the two adjacent CCPs as collocated, dividing workload, staffing, and resources equally between them (see Table 5).

Nurse supervisors are again a critical resource, and the CCPOM recommends allocating all available resources to the three CCPs. In contrast, ambulance seats are no longer a limiting constraint; of the 178 seats available, only 123 are needed.

	CCP			
Resource (units)	NM-3/4	NM-4/6	LS	Used resource (vs. available)
Decontamination units (units)	3	3	2	8 (25)
Triage EMS (persons)	15	15	12	42 (72)
Administrators (persons)	10	10	5	25 (300)
EMS nurse supervisors (persons)	60	60	30	150 (150)
Behavioral staff (persons)	3	3	3	9 (60)
Security, command and control (persons)	12	12	12	36 (112)
Transportation preparation (persons)	3	3	3	9 (30)
Ambulance spaces (seats)	45.6	45.1	32	122.7 (178)
Maximum throughput (casualties/hour)	250	250	125	

Table 5: For the CCPs, the table lists the resource utilization in the mall scenario. The last column compares the amount of resources allocated to all CCPs with the amount of resources originally available for each resource category.

The CCPOM predicts that only 11 shelters (of 29 available) will be needed to accommodate all displaced casualties. By coincidence, the number of chosen shelters is identical to the number chosen in the baseline scenario. The CCPOM solution expresses a preference for shelters with available capacity closer to the CCPs; it fills all but two of those shelters to capacity.

The same basic premises apply for evacuees. Accommodating all 6,500 evacuees requires 15 shelters. By comparison, the baseline scenario requires 16 shelters.

Regarding transportation to hospitals, the CCPOM distributes the 1,300 casualties requiring hospitalization among all 22 hospitals in the National Capital Region that report surge capacity. Unlike the baseline scenario, INOVA Fairfax Hospital is included, although at a relatively low utilization level.

The overall primary objective for this scenario is $z_1 = 1{,}381$ weighted casualties per hour, insignificantly less than in the baseline scenario (1,390 casualties per hour). However, for the secondary objective, $z_2 = 10{,}124$ hours is significantly worse than in the baseline scenario (8,242 hours, an increase of 22.8 percent). In this scenario, for casualties traveling from CCPs to shelters, the average transit time increases by almost 20 minutes per casualty.

One-Mile Scenario

The second rule modeled for comparison purposes mandates selecting CCPs within one statute mile of the IS. Only 11 of 52 total candidate CCPs qualify; from these 11, the CCPOM chooses five CCPs as optimal, all of which differ from the CCPs selected in the two previous scenarios. As in the baseline and mall scenarios, all 13,000 casualties are treated at the CCPs (see Table 6).

CCP	First wave	Second wave	Total
MCI Center (MCI)	391	1,229	1,620
Walker Jones Education Center (WJ)	532	1,674	2,206
Ludlow Taylor Elementary School (LT)	916	3,404	4,320
McKinley Technical High School (MT)	852	2,708	3,560
U.S. Capitol—South Lawn (CSL)	309	985	1,294
Total	3,000	10,000	13,000

Table 6: For each wave of arrivals, the table lists the CCPs selected in the one-mile scenario and the distribution of casualties that each services. The CCPOM projects all casualties can be treated and sent to their final destinations within the 48-hour planning horizon.

Resource (units)	CCP MCI	WJ	LT	MT	CSL	Used resource (vs. available)
Decontamination units (units)	1	2	3	2	1	9 (25)
Triage EMS (persons)	12	12	12.4	12	12	60.4 (72)
Administrators (persons)	3.1	4.3	8.3	6.9	3	25.6 (300)
EMS nurse supervisors (persons)	18.7	25.5	49.5	41.3	15	150 (150)
Behavioral staff (persons)	3	3	3	3	3	15 (60)
Security, command and control (persons)	12	12	12	12	12	60 (112)
Transportation preparation (persons)	3	3	3	3	3	15 (30)
Ambulance spaces (seats)	37.3	26	62.5	40.7	11.5	178 (178)
Maximum throughput (casualties/hour)	77.5	106.3	206.4	172.2	62.6	

Table 7: For the CCPs, the table lists the resource utilization in the one-mile scenario. The last column compares the amount of resources that are allocated to all CCPs with the amount of resources that were originally available for each resource category.

Results for resource utilization are analogous to the results from the previous scenarios. Nurse supervisors continue to be a critical resource. Ambulance seats are fully utilized, as they were in the baseline scenario. With the exception of those two categories, the model does not fully allocate any resources (see Table 7).

Shelters and their expected number of casualties differ from those in previous scenarios. However, in each scenario, the CCPOM shows that ample systemwide shelter capacity for both short- and long-term needs is available.

The model details the movement of casualties requiring hospitalization from the CCPs to area hospitals. Trends noted in the first two scenarios continue for the incumbent scenario. All 22 hospitals receive casualties. However, the proportion of total casualties received at each hospital changes from scenario to scenario; this reflects the different distances from CCPs in the three scenarios studied.

Specifically, in the one-mile scenario, we observe shorter trips to shelters: the secondary objective improves to $z_2 = 7{,}775$ hours (six percent better than in the baseline scenario—on average, 4.3 minutes per casualty). This occurs at the expense of worsening the primary objective: $z_1 = 1{,}317$ weighted casualties per hour, a 5.5 percent decrease (see Table 8).

Scenario	z_1 (maximize weighted casualties/hour)	z_2 (minimize traveling hours)
Baseline	1,390	8,242
Mall	1,381	10,124
One mile	1,317	7,775

Table 8: The table shows an objective function comparison of the three scenarios.

By interacting with the District planners, which led to the development of the two additional scenarios, we realized that the three scenarios offered three different CCP configurations, yet similar throughput, as measured by z_1. Our recommendation to add a secondary objective in the form of travel time to shelters supplemented the planners' understanding of their system. Specifically, they realized the significance of previously unidentified derivative effects, such as increased travel time.

Use of Results by District Planners

District planners deem many nonintuitive CCPOM results to be significant to their planning, programming, and budgeting efforts. For planning purposes, the model's categorized resource utilization was one of the most important contributions. For example, prior to the study, anecdotal evidence suggested the need to acquire large numbers of decontamination units, with commensurate increases in trained decontamination personnel. Contrary to this previously unverified need, the model consistently predicted an excess capacity in decontamination units. Consequently, the District was able to considerably reduce the number of such units in its planning assumptions, and decrease or eliminate programmed purchases of decontamination equipment. As a second-order effect, the substantial cost savings achieved allowed the District to use these funds to help it secure other resources that were originally underestimated. For example, based on persistent shortfalls revealed by the study, the need estimates for nurse supervisors and ambulances were modified. Thus, apart from selecting CCPs, the model has been instrumental in providing the planners with a verifiable decision support tool for resource allocation, with key insights into the planning, programming, and budgeting processes.

In addition to estimating resource requirements, the model yielded other nonintuitive results. Conventional wisdom might suggest that optimal throughput could be achieved by activating CCPs in order, from the closest to the incident site to the most distant. The model results, however, showed that the optimal order and timing of CCP activation was not necessarily linear. In this way, the model encouraged planners to challenge untested assumptions, and evaluate previously unconsidered alternatives. For example, they had not considered locating two CCPs side by side. Quite unexpectedly, both the baseline and mall scenarios recommended collocating two CCPs at the National Mall. District plans were updated based on this result; the planners now consider this alternative an integral part of planning for both national special security event planning and everyday events.

This study and its results have prompted further inquiry by the District; specifically, these results have become important inputs for a larger study, *National Capitol Region, Key Response Planning Factors for the Aftermath of Nuclear Terrorism*, which Lawrence Livermore Labs is doing for the U.S. Department of Homeland Security.

Extension to Multiple Incident Sites

To understand the limitations of the CCP model capabilities to handle more complex scenarios, we solved a number of instances, including different origins for the casualties. These could be used to model multiple incident sites, or only one IS, which affects a large area that we wish to divide.

Specifically, based on our baseline scenario, we investigated adding three new incident sites, populated with first-wave casualties who are subtracted from either the initial IS or from the second wave. In some instances, the new problems require up to three hours to be solved within one percent from optimal. This computational time may still be considered acceptable for planning purposes.

The results show that, on average, approximately 50 percent of the selected CCP locations change with respect to those in the baseline case. The distribution of the resources across CCPs is similar in all cases; most importantly, the critical limiting resources creating the bottlenecks (maximum throughput per hour at each CCP) are the same as in the single-IS case: EMS nurse supervisors and ambulance spaces.

Conclusions

CCPOM is an analytical engine to support strategic planning for future incidents. District planners have validated it for a stated, static set of conditions. We developed and refined the model through an iterative, collaborative process that effectively merges the academic realms of optimization modeling and management science with the operational experience and institutional knowledge of District planners and emergency responders.

The model generates discrete combination of CCPs, staffing levels, and ambulance allocations that satisfy the primary objective of generating the highest total-victim throughput system wide. Collaterally, the CCPOM also optimizes the distribution of casualties from both waves to the CCPs, and the movement of casualties from the CCPs to shelters and hospitals.

Although we investigated different scenarios to provide emergency planners with nonintuitive insights into persistent resource overages and shortfalls, our study is limited in scope to the conditions provided by the District. Its conclusions are relevant for a given range of conditions. However, District planners are faced with the challenge of planning for a myriad of scenarios, affecting parts of the region differently and requiring different responses, all in an environment of constrained resources and constricting budgets. The CCPOM has sufficient flexibility to allow it to be used for operational testing and further what-if analysis; thus, District planners can expand the model's utility by identifying additional areas for investigation. For example, future excursions of the model may include extensive sensitivity analysis for resource allocation in case of multiple incident sites (for which the model is already established), and with alternative objectives and uncertainties.

Appendix

This appendix describes the mathematical formulation for the CCPOM.

Sets:

$i \in I$, $c \in C$, $h \in H$, $l \in L$	incident site(s), candidate CCPs, hospitals, and shelters, respectively.
$r \in R$	Resources, $R = R^D \cup R^C$, where R^D is the subset of discrete resources (e.g., decontamination units) and R^C is the subset of continuous resources (e.g., nurses).
$t \in T$	Periods $\{0, 1, \ldots\}$.

Parameters [units]:

d	Duration of each period (0.25 hours in our scenario) [hours].
k_i^I	Casualties in a first wave, originating from IS i [casualties].
k^W	Casualties in a second wave, from unspecified origins [casualties].
t_c^{setup}	Number of periods to set up CCP c before treating any casualty [# periods].
f_{rt}^R	Available fraction of resource r during period t [fraction].
p_{ict}^I	Fraction of those arriving at CCP c from IS i that would do so at the beginning of period t [fraction].
p_t^W	Fraction of those arriving at all CCPs from unspecified origins at the beginning of period t [fraction].
	(Note: Both p_{ict}^I and p_t^W are derived from a given deterministic fraction of casualties arriving at CCPs; see Figures 1 and 2, respectively.)
t_{cl}^L, t_{ch}^H	Average time for casualties to travel from CCP c to shelter s (using public transportation or by foot) and to hospital h (by ambulance), respectively [hours].
q_c^C, q_l^L, q_h^H	Queuing capacity at CCP c, and physical capacity at shelter l and hospital h [persons].
q^B	Total ambulance capacity available (including ambulance bus) [seats].
f^{LLT}, f^{LST}, f^H	Fraction of casualties who need to go to shelters (long term), shelters (short term), and hospitals, respectively, after being treated at the CCP [fraction].
a_r^0, u_r	Initially available amount of resource r [resource-unit hours per hour] and processing capacity per unit of resource r [casualties/resource-unit hour], respectively.
$\underline{m}_r$, $\bar{m}_r$	Minimum and maximum of resource r at any CCP, respectively [resource units].
$\underline{n}$, $\bar{n}$	Minimum and maximum number of CCPs to locate, respectively [# of CCPs].
λ	Small penalty to ensure no unnecessary resources are allocated; we set this penalty to 0.0001 in all our instances.

Decision variables:

Y_c	1 if candidate CCP c is used as a CCP, and 0 otherwise.
F_{ic}^I, F_c^W, F_{cl}^{LLT}, F_{cl}^{LST}, F_{cht}^H	Flow of casualties from IS i to CCP c, from unspecified origin to CCP c, from CCP c to shelter l (for long-term stay), from CCP c to shelter l (for short-term stay), and from CCP c to hospital h in period t, respectively [casualties].
Q_{ct}, A_{ct}, S_{ct}	Queue length at the end of, arrivals at the beginning of, and casualties served (treated) during period t, respectively, at CCP c [casualties].

X_{cr}	Number of resource r allocated to CCP c [resource-unit hours per hour].
B_c	Number of ambulance seats allocated to CCP c [seats].
z_1, z_1^*	Primary objective function, and its optimal value, respectively [weighted casualties].
z_2, z_2^*	Secondary objective function, and its optimal value, respectively [hours].

Formulation:

$$\sum_c F^I_{ic} = k^I_i \quad \forall i. \tag{1}$$

$$\sum_c F^W_c = k^W. \tag{2}$$

$$A_{ct} = \sum_i p_{ict} F^I_{ic} + p^W_t F^W_c \quad \forall c, \forall t > 0. \tag{3}$$

$$Q_{c0} = A_{c0} = S_{c0} = 0 \quad \forall c. \tag{4}$$

$$Q_{ct} = Q_{c,t-1} + A_{ct} - S_{ct} \quad \forall c, \forall t > 0. \tag{5}$$

$$Q_{ct} \le q^C_c Y_c \quad \forall c, \forall t > 0. \tag{6}$$

$$S_{ct} = 0, \quad \forall c, \forall 0 < t < t^{\text{setup}}_c. \tag{7}$$

$$S_{ct} \le d u_r f^R_{rt} X_{cr} \quad \forall c, \forall t > 0, \forall r \in R. \tag{8}$$

$$\sum_c X_{cr} \le a^0_r \quad \forall r \in R. \tag{9}$$

$$\underline{m}_r Y_c \le X_{cr} \le \bar{m}_r Y_c \quad \forall c, r \in R. \tag{10}$$

$$\sum_l F^{LLT}_{cl} = \sum_t f^{LST} S_{ct} \quad \forall c. \tag{11}$$

$$\sum_l F^{LST}_{cl} = \sum_t f^{LST} S_{ct} \quad \forall c. \tag{12}$$

$$\sum_c F^{LLT}_{cl} \le q^{LST}_l \quad \forall l. \tag{13}$$

$$\sum_c F^{LST}_{cl} \le q^{LST}_l \quad \forall l. \tag{14}$$

$$\sum_h F^H_{cht} = f^H S_{c,t-1} \quad \forall c, \forall t. \tag{15}$$

$$\sum_h 2 t^H_{ch} F^H_{cht} \le d B_c \quad \forall c, \forall t. \tag{16}$$

$$\sum_c \sum_t F^H_{cht} \le q^H_h \quad \forall h. \tag{17}$$

$$\sum_c B_c \le q^B. \tag{18}$$

$$\underline{n} \le \sum_c Y_c \le \bar{n}. \tag{19}$$

$$F^I_{ic}, F^W_t, Q_{ct}, A_{ct}, S_{ct}, F^{LLT}_{cl}, F^{LST}_{cl}, F^H_{cht}, B_c \ge 0 \quad \forall i, \forall c, \forall l, \forall h, \forall t. \tag{20}$$

$$X_{cr} \ge 0 \quad \forall c \in C, \; r \in R^C. \tag{21}$$

$$X_{cr} \ge 0 \quad \text{and integer } \forall c \in C, \; r \in R^D. \tag{22}$$

$$Y_c \in \{0, 1\} \quad \forall c \in C. \tag{23}$$

$$z_1^* = \max z_1, \quad \text{where } z_1 = \sum_c \sum_{t>0} d^{-1} t^{-1} S_{ct}. \tag{24}$$

$$z_2^* = \min z_2, \quad \text{where } z_2 = \sum_c \sum_l t^L_{cl} (F^{LLT}_{cl} + F^{LST}_{cl}) + \lambda \sum_c \sum_r X_{cr}. \tag{25}$$

$$z_1 \ge 0.99 z_1^*. \tag{26}$$

CCPOM-1: $z_1^* = \max z_1$, subject to (1)–(23).
CCPOM-2: $z_2^* = \min z_2$, subject to (1)–(23), (26).

Description of the formulation:

Constraint (1) ensures all casualties from each IS are distributed across the CCPs. Constraint (2) does the same, but for second-wave casualties (from unspecified origins). In both cases, we assume emergency staff members can direct individual and (or) groups of casualties to whichever CCPs they deem appropriate. Constraint (3) assumes that the arrival of casualties at any given CCP follows a prespecified distribution over time. Typically, the distribution is based on the distance between the IS and the CCP: The farther away they are from each other, the more dispersed the arrival distribution will be. This distribution applies to the total number of casualties the model decides should be directed to travel from each IS to each CCP. The same concept applies to second-wave casualties. Constraints (4)–(6) handle queues at the CCPs: Constraint (4) establishes the initial conditions at time 0. Constraint (5) ensures the queue length increases with casualty arrivals and decreases with treated casualties. Constraint (6) ensures queues (at established CCPs) never exceed the CCP capacities. Constraint (7) ensures that no treatment is provided before the CCP has been set up (even if arrivals have occurred), and constraint (8) limits that service based on the resources allocated to the CCP. Note that this occurs for each resource individually considered; thus, the most restrictive one applies in each period. Constraints (9) and (10) distribute available resources within the limits of each CCP. Constraints (11)–(14) ensure the appropriate fractions of treated casualties are distributed across long- and short-term shelters, without exceeding their capacities. No specific constraint is associated with the timing in this process because these patients may wait at the CCP for public transportation (or walk to the shelter). Traveled distance is still penalized in the second objective function (see next). However, constraints (15)–(18) ensure that sufficient ambulance seats are allocated to the CCP on a period-to-period basis. This takes into account: (i) the fraction of casualties that will be transported to hospitals; (ii) assumed two-way trips by the ambulances; (iii) hospital capacity; and (iv) total ambulance seats available. Constraint (19) limits the number of CCPs selected. Constraints (20)–(23) establish the domain for decision variables.

Objective function (24) maximizes weighted throughput of casualties, where weights are inversely proportional to the time when casualties are treated to give higher priority to casualties treated early. This constitutes our first model, CCPOM-1. Constraint (25) is a secondary goal, optimized after objective function (24), which minimizes travel time to shelters for all casualties in the system, but subject to constraint (26), which establishes the threshold on the primary goal. This model, CCPOM-2, is executed after CCPOM-1.

References

Balcik B, Beamon BM (2008) Facility location in humanitarian relief. *Internat. J. Logist. Res. Appl.* 11(2):101–121.

Brandeau ML, Chiu SS (1989) An overview of representative problems in location research. *Management Sci.* 35(6):645–674.

Caunhye A, Nie X, Pokharel S (2012) Optimization models in emergency logistics: A literature review. *Socio-Econom. Planning Sci.* 46(1):4–13.

Church R, ReVelle CR (1974) The maximal covering location problem. *Papers Regional Sci.* 32(1):101–118.

Daskin MS (1995) *Network and Discrete Location: Models, Algorithms and Applications* (John Wiley & Sons, New York).

Daskin MS, Dean LK (2004) Location of health care facilities. Sainfort F, Brandeau M, Pierskalla W, eds. *Handbook of OR/MS in Health Care: A Handbook of Methods and Applications* (Kluwer, Norwell, MA), 43–46.

Dekle J, Lavieri MS, Martin E, Emir-Farinas H, Francis RL (2005) A Florida county locates disaster recovery centers. *Interfaces* 35(2):133–139.

District of Columbia Hospital Association (2006) 2006 hospital capability assessment for readiness. Report, District of Columbia Hospital Association, Washington, DC.

Drezner T (2004) Location of casualty collection points. *Environ. Planning C: Government Policy* 22(6):899–912.

Drezner T, Drezner Z, Salhi S (2005) A multi-objective heuristic approach for the casualty collection points location problem. *J. Oper. Res. Soc.* 57(6):727–734.

Eiselt HA, Laporte G (1995) Objectives in location problems. Drezner Z, ed. *Facility Location: A Survey of Applications and Methods* (Springer-Verlag, New York), 151–180.

Erkut E (1993) Inequality measures for location problems. *Location Sci.* 1(3):199–217.

Federal Emergency Management Agency (2009) FEMA fact sheet: National planning scenarios. Accessed May 1, 2014, http://www.hsdl.org/?view&did=683091.

Federal Emergency Management Agency (2010) Planning guidance for a response to a nuclear detonation, 2nd ed. Accessed May 1, 2014, http://www.epa.gov/radiation/docs/er/planning-guidance-for-response-to-nuclear-detonation-2-edition-final.pdf.

GAMS Development Corporation (2013) *GAMS—A User's Guide* (GAMS Development Corporation, Washington, DC).

Google (2012) Google Maps for business: Google Earth Pro. Accessed May 1, 2014, http://www.google.com/enterprise/earthmaps/earthpro.html.

Hakimi SL (1964) Optimal locations of switching centers and the absolute centers and medians of a graph. *Oper. Res.* 12(3):450–459.

Hakimi SL (1965) Optimum distribution of switching centers in a communication network and some related graph theoretic problems. *Oper. Res.* 13(3):462–475.

Lee E, Smalley H, Zhang Y, Pietz F, Benecke B (2009) Facility location and multi-modality mass dispensing strategies and emergency response for biodefense and infectious disease outbreaks. *Internat. J. Risk Assessment Management—Biosecurity Assurance Threatening World: Challenges Explorations Breakthroughs* 12(2–4):311–351.

Maimon O (1986) The variance equity measure in locational decision theory. *Ann. Oper. Res.* 6(5):147–160.

Marianov V, ReVelle C (1995) Siting emergency services. Drezner Z, ed. *Facility Location: A Survey of Applications and Methods* (Springer-Verlag, New York), 199–223.

Owen S, Daskin M (1998) Strategic facility location: A review. *Eur. J. Oper. Res.* 111(3):423–447.

Verification Letter

John Donnelly, Battalion Chief, Homeland Security, Office of Domestic Preparedness, District of Columbia Fire and EMS, 1338 Park Road NW, Washington, DC 20010, writes:

"Aruna, Curtis, and Javier,

"I have just finished reviewing the article, 'Casualty Collection Points Optimization: A Study for the District of Columbia.' This article brings back fond memories of working on this project and the team we had assembled to address this problem.

"The article accurately portrays the process and the broad scale of participation required to identify and collect the data used by the Casualty Collection Points Optimization Model. The results of the planning and modeling effort were used to update and inform and District Response Plan and the District of Columbia Fire and Emergency Medical Services Department's Mass Decontamination Plan.

"I endorse your team's article and support its publication."

Aruna Apte received her PhD in operations research from Southern Methodist University in Dallas and her master's in mathematics from Temple University in Philadelphia. She is an associate professor in operations and logistics management in the Graduate School of Business and Public Policy at the Naval Postgraduate School. She teaches mathematical modeling, for which she won the best teacher award. She has advised more than 60 students for MBA/master's reports. Her research interests are in developing mathematical models for complex, real-world operational problems using optimization tools. Her research is focused on humanitarian and military logistics. Before NPS, she worked as a consultant at MCI and taught at the Cox School of Business, SMU, where she won the best teacher award. She served as the president for Humanitarian Operations and Crisis Management College in Production and Operations Management Society.

Curtis Heidtke is an attorney-advisor at the U.S. Naval Postgraduate School in Monterey, California. At the time of the execution of this project he was a homeland security consultant based in Washington, DC, providing federal, state, and local governments and private entities with strategic solutions to homeland security-related problems. He was formerly with the Office of General Counsel, Intelligence and Analysis Division at the U.S. Department of Homeland Security. Heidtke is a retired U.S. Air Force judge advocate. He is a command pilot with over 2,500 hours in high-performance military aircraft. He holds a Bachelor of Science from the U.S. Air Force Academy, a Master of Science in systems management from the University of Southern California, a Juris Doctorate from the University of Michigan, and a Master of Arts in Homeland Security and Defense from the Naval Postgraduate School.

Javier Salmerón received his MS and PhD in mathematics from Complutense University of Madrid and Polytechnic University of Madrid, respectively. He is an associate professor in the Operations Research Department at the Naval Postgraduate School (NPS), where he teaches courses and performs research in optimization and its applications. Before his time with NPS, Dr. Salmerón was with the engineering branch of Spanish electric utility Iberdrola and an adjunct professor in the Department of Statistics and Operations Research at the Statistics School of Complutense University.