OPERATIONS RESEARCH
Vol. 56, No. 1, January–February 2008, pp. 247–254
ISSN 0030-364X | EISSN 1526-5463 | 08 | 5601 | 0247

DOI 10.1287/opre.1070.0490


TECHNICAL NOTE

# Spatial Queueing Analysis of an Interdiction System to Protect Cities from a Nuclear Terrorist Attack

**Michael P. Atkinson**
Institute for Computational and Mathematical Engineering, Stanford University, Stanford, California 94305, mpa33@stanford.edu

**Lawrence M. Wein**
Graduate School of Business, Stanford University, Stanford, California 94305, lwein@stanford.edu

We formulate and analyze a spatial queueing model concerning a terrorist who is attempting to drive a nuclear or radiological weapon toward a target in a city center. In our model, imperfect radiation sensors form a circular wall around the periphery of the city, and vehicles setting off sensor alarms (representing a terrorist or a nuisance alarm) arrive randomly at the perimeter of a circle (representing the wall of sensors) and drive toward the center of the circle. Interdiction vehicles, one in each wedge of the circle, chase the alarm-generating vehicles. We derive an accurate mathematical expression for the mean damage inflicted by a terrorist in this system in terms of the arrival rate of alarm-generating vehicles and the number of interdiction vehicles. Our results suggest that detection-interdiction systems using current technology are capable of mitigating the damage from a nuclear weapon made of plutonium, but not one made of uranium or a radiological weapon.

*Subject classifications*: government: defense; queues: approximations.
*Area of review*: Military and Homeland Security.
*History*: Received November 2005; revisions received October 2006, May 2007, August 2007; accepted August 2007.

## 1. Introduction

A nuclear weapon (made of uranium or plutonium) detonated by a terrorist in a large U.S. city could kill a half-million people and cause one trillion dollars in direct economic damage (Bunn et al. 2003). Although this threat is deemed "real and urgent" (Bunn et al. 2003, p. viii), a more likely scenario is for terrorists to assemble a radiological dispersal device (or so-called dirty bomb) containing radiological material such as cesium, which would inflict much less damage, but would nonetheless wreak considerable havoc. Because the majority of nuclear material in the former Soviet Union remains vulnerable to theft (Bunn et al. 2003), and smuggling nuclear or radiological material into a U.S. port in a shipping container is fairly easy (Flynn 2000, Stana 2004), we focus on our last line of defense, which is to detect an assembled nuclear or radiological weapon as it is driven into a city and to provide timely and effective interdiction. Indeed, the U.S. government is in the process of developing (U.S. Dept. of Homeland Security 2005) and deploying (including a pilot test in New York City; Lipton 2007) such detection-interdiction systems in its largest cities, and our goal is to perform a rough-cut feasibility analysis. In this paper, we focus on the interdiction aspects of these systems by formulating and analyzing a spatial queueing model that incorporates scarce interdiction resources and implicitly accounts for false positives. As explained in §5, we embed our results into a Stackelberg game in a companion paper (Wein and Atkinson 2007) to analyze the entire detection-interdiction system.

Our interdiction model, which is formulated in §2, considers a spatial queueing system in a circle, whose perimeter represents the wall of radiation sensors. Customers correspond to vehicles triggering a sensor (typically a false alarm), and servers are interdiction vehicles. Customers arrive randomly on the perimeter and drive directly toward the circle center, and an interdiction vehicle chases each customer. Following Bertsimas and van Ryzin (1993), who look at a related mobile-server model in which customers arrive randomly in the circle but are immobile, we divide the circle into wedges and place an interdiction vehicle in each wedge, with a resting location that minimizes the mean customer sojourn time. We derive an analytically tractable expression in §3 that accurately approximates (because we have been unable to analytically determine the accuracy of these approximations, we assess their accuracy via simulation in §4) the primary performance measure of the queueing system, which is the mean damage caused by a terrorist. Concluding remarks are provided in §5, including a brief discussion of radiation emissions and detection (which is needed to understand the implications of our results, but not to understand the model and analysis) and the technological and logistical feasibility of detection-interdiction systems.

Network interdiction is an active field of study within operations research (e.g., Woodruff 2003). Many of the early problems considered maximizing an adversary's shortest path, or minimizing an adversary's maximum flow, through a network. Several authors (Wollmer 1964, Washburn and Wood 1995, Pan et al. 2003), motivated by military operations, or smuggling of nuclear materials or drugs, allow the inspector to locate detectors at certain arcs and permit the adversary to choose a path through the network to maximize his probability of evading detection. Relative to these papers, our model makes the simplifying assumption that the detection probability of sensors is independent of location. While this assumption seems reasonable in our context because the same technology is used throughout the system, it may be violated if vehicle speeds at different highway ramps (dictated by the ramp curvature) and/or length of red lights at traffic intersections differ appreciably. Our grid is also more restrictive than the general networks considered in the literature. However, our model is more complex than those in the literature in that it takes a more detailed approach to the interdiction process; to our knowledge, this study contains the first queueing analysis of a system with mobile servers and mobile customers.

## 2. Model Formulation

We consider a circle of radius $R$, where the target is in the middle of the circle and the perimeter of the circle represents an outer radiation wall. The goal of our analysis is to compute the mean damage caused by a terrorist. However, because a terrorist with a nuclear weapon is an extremely rare event, the congestion in our queueing model is dictated by false positive alarms. That is, the customers in our model correspond to vehicles that set off a nuisance alarm when passing through the wall of sensors. Customers arrive according to a Poisson process at rate $\lambda$ and appear uniformly on the circle perimeter. We give the position of the customer and server in polar coordinates, with the target being the origin. A customer arriving at position $(R, \theta_c)$ immediately starts traveling at velocity $R$ per hour in a straight line toward the circle center (so that it takes one hour to reach the center). We denote his generic location in the system as $(r_c, \theta_c)$.

The $M$ servers in the model are interdiction vehicles. Following Bertsimas and van Ryzin (1993), we divide the circle into $M$ equal-sized wedges and assign one vehicle to service the customers in each wedge, thereby reducing the analysis to a single-server queueing system on a wedge. The interdiction vehicles travel at velocity $\alpha R$, where $\alpha > 1$. To mimic the gridlike nature of roads, interdiction vehicles are restricted to moving in a polar manner along rays (i.e., in a direction emanating out from the circle) and arcs (i.e., constant-radius paths), as described, e.g., in Larson and Odoni (1981, p. 175). We also analyzed the case in which interdiction vehicles are not restricted in their paths (analysis not shown), and derived qualitatively similar results. Although allowing vehicles to travel to adjacent wedges is not apt to improve performance very much in light traffic (see §4), restricting each vehicle to a wedge (as we have done) does make the system vulnerable to a multipronged attack in which two vehicles (with the first being a decoy) arrive nearly simultaneously to the same wedge.

Because the practically relevant regime for this problem is light traffic, our queueing discipline mimics a policy shown to minimize the mean customer sojourn time in light traffic in a related system (Bertimas and van Ryzin 1993). More specifically, we assume that when there are no customers in the system, the interdiction vehicle is located at the optimal resting location, which minimizes the expected sojourn time of a customer arriving to an empty system; this location is derived later in this section. If a customer arrives to an empty system, then the server travels on an interdiction path toward the customer so as to minimize the distance the customer travels (this is computed in Proposition 1 below). Upon catching the customer, the interdiction vehicle performs an on-site service that is a random variable with mean $m_s$ and standard deviation $\sigma_s$; our approximation scheme uses only the first two moments of the on-site service time, although we assume this service time is normally distributed in the computational study in §4. When the on-site service is completed, the customer instantaneously exits the system, and the interdiction vehicle travels back to the optimal resting location if there are no other customers in the system. If there are other customers in the system, they are immediately (i.e., without the server returning to the optimal resting location) served in a first-come first-served manner, with one caveat: the interdiction vehicle first computes whether or not the customer can be caught before reaching the target. If the customer is catchable (i.e., if $r_c > r_s/\alpha$, where $r_c$ and $r_s$ are the radii of the customer and server at the service completion epoch of the previous customer), then the interdiction vehicle chases the customer in a time-minimizing manner; if the customer is not catchable, then the interdiction vehicle does not pursue him. Also, if a new customer arrives as the interdiction vehicle is traveling back to its resting location, the interdiction vehicle immediately starts chasing the new customer.

Without loss of generality, we assume that the wedge spans the angles $[0, 2\pi/M]$. The basic building block of our model is the computation of the distance a customer arriving at generic location $(r_c, \theta_c)$ travels before being caught by an interdiction vehicle starting from the generic location $(r_s, \theta_s)$. Proposition 1 below is derived by comparing the arc/ray strategy, in which the interdiction vehicle catches the customer by moving first along the arc from $\theta_s$ to $\theta_c$ and then along the ray generated by $\theta_c$, to the ray/arc strategy, in which the interdiction vehicle moves along the ray generated by $\theta_s$ and then along the arc from $\theta_s$ to $\theta_c$.

It turns out that ray/arc strategy is optimal if the angular difference between the server location and customer location is large (the first case in Equation (1)) or if the interdiction radius (i.e., the radial location where interdiction occurs) is less than $r_s$ (the second case in (1)); and the arc/ray strategy is optimal if the interdiction radius is greater than $r_s$ (the third case in (1)). We show in the proof of Proposition 1 (in the online companion) that the distance defined in (1) is the minimum distance a customer can travel when the interdiction vehicle's chase path is restricted to rays and arcs. An electronic companion to this paper is available as part of the online version that can be found at http://or.pubs.informs.org/ecompanion.html.

Proposition 1. *Let $\psi = |\theta_c - \theta_s|$. Then, the distance a customer travels before interdiction when the customer starts at $(r_c, \theta_c)$ and the interdiction vehicle starts at $(r_s, \theta_s)$ is*

$$d(r_c, r_s, \psi) = \begin{cases} \dfrac{r_c + r_s}{\alpha + 1} & \text{if } 2 \leqslant \psi, \\[2ex] \dfrac{r_s - r_c + r_c\psi}{\alpha - 1 + \psi} & \text{if } \dfrac{\alpha(r_c - r_s)}{r_s} \leqslant \psi < 2, \\[2ex] \dfrac{r_c - r_s + r_s\psi}{\alpha + 1} & \text{if } \psi < \dfrac{\alpha(r_c - r_s)}{r_s}. \end{cases} \quad (1)$$

By uniformity and symmetry, the optimal resting location is $(r^*, \pi/M)$, where $r^*$ is the optimal radius. To determine $r^*$, we use (1) and numerically compute

$$r^* = \arg\min_r \int_0^{\pi/M} d(R, r, \psi)\, d\psi. \quad (2)$$

Finally, to compute the expected damage, we consider a terrorist arriving in steady state. We assume that the damage function is $b - ar$ if a terrorist detonates a bomb at radius $r$, and we set $b = 10$ and $a = 9/R$ to maintain a 1-to-10 damage scale. If the terrorist is not catchable, then he detonates at $r = 0$. If he is caught at radius $r$, then he detonates the bomb at radius $r$ with probability $q$.

The use of a linear damage function is a gross simplification. There are four main effects of a nuclear weapon: shock and blast, thermal radiation, initial nuclear radiation, and residual nuclear radiation (Glasstone and Dolan 1977). All four exposures are nonlinear functions of distance (thermal radiation varies inversely with distance squared, and radiation exposure varies inversely with distance squared with scattering and decreases exponentially with no scattering) and depend greatly on the yield of the bomb. Moreover, the dose-response effects and the population gradient are also nonlinear. Although the yield of a bomb detonated by a terrorist is highly uncertain, the instantaneously fatal effects are on the order of miles (e.g., for the Hiroshima bomb) to tens of miles, and the residual effects (which could also eventually be fatal) are on the order of tens of miles. Although the linear damage function with a 1-to-10 scale allows policymakers to easily internalize our results (by interpreting damage as relative distance), we believe that our results would need to be combined with detailed simulation models of the four effects of a nuclear weapon (including dose-response models and spatial population data) to provide comprehensive input to policymakers; such an effort is beyond the scope of this paper.

## 3. Analytical Approximation

Because we require an analytically tractable version of the interdiction model to embed into an optimization framework in Wein and Atkinson (2007), and in an attempt to gain an understanding of how this system behaves as a function of $\lambda$ and $M$, in this section we approximate the mean damage caused by a terrorist.

Our analysis has three steps. The first step is to approximate the optimal resting location and the average distance a customer travels before being caught given that the server is idle and at the optimal resting location. We then approximate the spatial queue by an $M/M/1/2$ queue with reneging using the quantities derived in the first step. Finally, we derive the expected damage inflicted by a terrorist using the $M/M/1/2$ reneging model. The accuracy of each of the approximations (we use equal signs throughout) in this section is assessed in §4.

**Mean Travel Distance of Customers.** The calculation in Equation (2) is difficult because of the three-part expression for $d(R, r, \psi)$ in (1). We simplify the analysis derived in this first step by choosing our optimal resting location to be the maximum radius such that all interdictions take place at radii greater than or equal to the radius of the optimal resting location (i.e., all interdictions occur "in front" of the server). We only analyze $M \geqslant 2$, and thus we only need to consider the second and third cases in (1). We derive the value of $r$ by imposing equality in the second and third conditions on the right side of (1), i.e., $R - r\psi/\alpha = r$ for the maximum value of $\psi$, which is $\pi/M$, yielding the approximation

$$r^* = \frac{R}{1 + \pi/\alpha M}. \quad (3)$$

Using the $r^*$ in (3), we compute the mean travel distance of a customer during pursuit who enters an empty system with the server at the optimal resting location to be

$$\begin{aligned} E[d(R, r^*, \psi)] &= E\left[\frac{R - r^*(1-\psi)}{\alpha+1}\right] \quad \text{by (1)}, && (4) \\ &= \frac{R - r^*(1-\pi/2M)}{\alpha+1} \quad \text{because } E[\psi] = \frac{\pi}{2M}, \\ &= \frac{(\alpha+2)\pi R}{2(\alpha+1)(\alpha M+\pi)} \quad \text{by (3)}. && (5) \end{aligned}$$

**The $M/M/1/2$ Queue.** Motivated by the observation that for any realistic value of $\alpha$ there is almost no chance of catching a customer who arrives to find more than one other customer in the system, we use (5) to approximate the queueing system by an $M/M/1/2$ queue with arrival rate $\lambda/M$ and service rate $\mu$, which will be calculated shortly. The exponential service time approximation might perform well because the queueing system has behavior similar to an infinite-server queue (because it is in light traffic) and a loss system (because it has only one buffer space), both of which have performance measures that are insensitive to the service-time distribution (Gross and Harris 1985). In addition, we allow reneging of customers arriving to find one other customer in the system because they may be uncatchable. We assume that the time until reneging for customers that arrive to find one customer in the system is an exponential random variable with mean $r_1^{-1}$, so that the steady-state probability of arriving to this queue when it has zero, one, or two customers, respectively, is (Gross and Harris 1985)

$$p_0 = \frac{1}{1+\lambda/(M\mu)+\lambda^2/(M^2\mu(\mu+r_1))},$$
$$p_1 = \frac{\lambda/(M\mu)}{1+\lambda/(M\mu)+\lambda^2/(M^2\mu(\mu+r_1))}, \quad \text{and} \tag{6}$$
$$p_2 = \frac{\lambda^2/(M^2\mu(\mu+r_1))}{1+\lambda/(M\mu)+\lambda^2/(M^2\mu(\mu+r_1))}.$$

It remains to specify the service rate $\mu$ and the reneging rate $r_1$. Service consists of two components, chase and on-site, and the mean chase time varies according to whether the server is idle or busy when a customer arrives. Recalling that $m_s$ is the mean on-site service time, we let

$$\mu^{-1} = m_s + m_t, \tag{7}$$

where $m_t$ is the mean chase time, which we compute using a fixed-point approach. Let $P_n$ be the fraction of customers arriving to find one customer in the system who are eventually caught. In Equation (14), we compute the reneging probability $P_r$, which equals $1-P_n$. Let $t_e$ be the mean time it takes the server to catch a customer when the server is idling at his resting location when the customer arrives, and let $t_b$ be the mean time it takes to catch a customer (who does not renege) when the server is busy at the time of the customer arrival. Dividing the expected distance the customer travels, given in Equation (5), by his speed, $R$ miles per hour, implies that

$$t_e = \frac{(\alpha+2)\pi}{2(\alpha+1)(\alpha M+\pi)}, \tag{8}$$

and $t_b = d_b/R$, where $d_b$, which is the distance traveled by a nonreneging customer after the chase begins, is calculated in Equation (20) when we estimate the damage inflicted by a terrorist. Therefore, $m_t$ satisfies

$$m_t = \frac{p_0}{p_0+P_np_1}t_e + \frac{P_np_1}{p_0+P_np_1}t_b \tag{9}$$
$$= \frac{M}{M+P_n\lambda(m_s+m_t)}t_e + \frac{P_n\lambda(m_s+m_t)}{M+P_n\lambda(m_s+m_t)}t_b \quad \text{by (6) and (7).} \tag{10}$$

Solving (10) for $m_t$ yields

$$m_t = \Big[-(M+\lambda P_n m_s - \lambda P_n t_b) + \sqrt{(M+\lambda P_n m_s-\lambda P_n t_b)^2+4P_n\lambda(Mt_e+P_n\lambda m_s t_b)}\Big] \cdot (2P_n\lambda)^{-1}, \tag{11}$$

and substitution of this quantity into (7) gives the service rate $\mu$ for the $M/M/1/2$ system (in terms of $P_n$ and $t_b$, which are calculated later).

In this $M/M/1/2$ system, the reneging probability of a customer who arrives to find one other customer in the system is $r_1/(r_1+\mu)$. In determining a value for $r_1$, we will attempt to reincorporate some of the nonexponential features of the spatial model described in §2, where a customer arriving to find one other customer in the system reneges if the time it takes for him to become uncatchable is less than the residual service time of the customer currently in service. More specifically, we choose $r_1$ so that $r_1/(r_1+\mu) = P(T_1 < T_2)$, where the random variables $T_1$ and $T_2$ are approximate representations of, respectively, the time until an arriving customer is uncatchable and the residual service time of the customer currently in service. To determine $T_1$, we note that if the server is located at $r_s$ at the time of customer arrival, the time until the customer becomes uncatchable is $1 - r_s/\alpha R$ because at this time point his radial location is $r_s/\alpha$. For simplicity, we assume that the customer currently in service arrived to an empty system with the server at his optimal resting location, so that by (3) and (4) the server location is uniformly distributed with parameters $(\alpha R + r^*)/(\alpha+1)$ and $r^*$. Hence, we assume that $T_1$ is a $U[r_l, r_u]$ random variable with parameters

$$r_l = \frac{\alpha^2 R - r^*}{\alpha(\alpha+1)R}, \quad r_u = 1 - \frac{r^*}{\alpha R}. \tag{12}$$

Because the residual service time has a complicated distribution, we make the simplifying assumption that $T_2$ is exponentially distributed with a mean $m_r$ (and rate $\mu_r = m_r^{-1}$) equal to the mean residual service time, where the $k$th moment of the residual service time equals the $k+1^{st}$ moment of the service time divided by $k+1$ times the mean service time (Equation (5-47a) in Heyman and Sobel 1982). By Equations (3) and (4), the chase time is a linear function of $\psi$, and thus is uniformly distributed between $\pi/(\alpha M+\pi)(\alpha+1)$ and $\pi/(\alpha M+\pi)$, which yields

$$\begin{aligned} m_r &= \mu_r^{-1} \\ &= \left[\sigma_s^2 + \frac{1}{12}\left(\frac{\pi\alpha}{(\alpha+1)(\alpha M+\pi)}\right)^2 + \left(\frac{(\alpha+2)\pi}{2(\alpha+1)(\alpha M+\pi)} + m_s\right)^2\right] \\ &\quad \cdot \left[2\left(\frac{(\alpha+2)\pi}{2(\alpha+1)(\alpha M+\pi)} + m_s\right)\right]^{-1}. \end{aligned} \tag{13}$$

If we denote the reneging probability, $P(T_1 < T_2)$, by $P_r$, where $P_n = 1 - P_r$ was used to compute the service rate $\mu$, then

$$P_r = \frac{e^{-\mu_r r_l} - e^{-\mu_r r_u}}{\mu_r(r_u - r_l)}, \tag{14}$$

and equating (14) to $r_1/(r_1+\mu)$ gives

$$r_1 = \frac{P_r\mu}{1-P_r}. \tag{15}$$

In our computational study, $T_1$ and $T_2$ are of comparable magnitude and $P(T_1 < T_2) \approx 0.3$.

**Expected Damage.** Finally, to compute the expected damage, we assume that with probability $p_2$ customers are not catchable and cause damage $b$, with probability $p_0$ customers are caught at average radius $R-(\alpha+2)R\pi/2(\alpha+1)(\alpha M+\pi)$ by Equation (4) and successfully detonate with probability $q$, and with probability $p_1$ a server does not begin the chase until after a residual (travel plus on-site) service time. For this last group of customers, a fraction $r_1/(r_1+\mu)$ renege, and hence cause damage $b$, and the only remaining difficulty is to estimate the mean distance the nonreneging customers travel before being caught (at which point they detonate with probability $q$). In terms of our earlier notation, the mean amount of time a nonreneging customer travels before the server begins chasing him is $E[T_2 \mid T_2 < T_1]$, i.e., the mean residual service time conditioned on it being less than the time until the customer is uncatchable. Assuming $T_2$ is exponential with the parameter in (13) leads to a significant underestimation of $E[T_2 \mid T_2 < T_1]$, perhaps because the coefficient of variation of $T_2$ is significantly less than one in our numerical computations. Consequently, we instead assume that $T_2$ is normal with mean $m_r$ in (13) and variance $\sigma_r^2$, which is the variance of the residual lifetime of the sum of the chase time ($U[\pi/(\alpha M+\pi)(\alpha+1), \pi/(\alpha M+\pi)]$) and the on-site service time, given by

$$\sigma_r^2 = \frac{1}{3}\left(\frac{(\alpha+2)\pi}{2(\alpha+1)(\alpha M+\pi)} + m_s\right)^2 + \sigma_s^2 + \frac{1}{12}\left(\frac{\pi\alpha}{(\alpha+1)(\alpha M+\pi)}\right)^2 - m_r^2. \tag{16}$$

Although computing $E[T_2 \mid T_2 < T_1]$ analytically is possible if $T_1$ is uniformly distributed with parameters in (12) and $T_2$ is normal, the resulting expression is tedious, and instead we replace $T_1$ by its mean, $(r_l+r_u)/2$. For the parameters used in this paper, the difference between assuming $T_1$ is uniform and $T_1$ is $(r_l+r_u)/2$ has a negligible impact on the computation of $E[T_2 \mid T_2 < T_1]$. With these assumptions, the nonreneging customer's mean radial location at the time the server starts chasing him is

$$\tilde{r}_c = R\left(1 - E\left[T_2 \mid T_2 < \frac{r_l+r_u}{2}\right]\right) \tag{17}$$

$$\begin{aligned} &= R\left\{1 - \left[\Phi\left(\frac{(r_l+r_u)/2 - m_r}{\sigma_r}\right)\right]^{-1} \cdot \left[m_r\Phi\left(\frac{(r_l+r_u)/2-m_r}{\sigma_r}\right) \right.\right. \\ &\qquad \left.\left. - \sigma_r\phi\left(\frac{(r_l+r_u)/2-m_r}{\sigma_r}\right)\right]\right\}, \end{aligned} \tag{18}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the probability density function and cumulative distribution function of the standard normal, respectively.

If we assume that the customer currently in service arrived while the server was in his optimal resting location, then the server's mean radial location is $\tilde{r}_s = R(1-(\alpha+2)\pi/2(\alpha+1)(\alpha M+\pi))$ when the chase begins. Although the server could be located anywhere along this arc of radius $\tilde{r}_s$, for analytic tractability we assume that the server is located in the middle of the arc at $\theta_s = \pi/M$. Using Proposition 1, the mean distance traveled after the chase starts, denoted by $d_b$, is

$$d_b = \frac{M}{\pi}\left(\int_0^{(\alpha/r_s)(r_c-r_s)} \frac{r_c - r_s + r_s\psi}{\alpha+1}\,d\psi + \int_{(\alpha/r_s)(r_c-r_s)}^{\pi/M} \frac{r_s - r_c + r_c\psi}{\alpha-1+\psi}\,d\psi\right). \tag{19}$$

Because we analyze $M \geqslant 2$, we do not need to look at the $\psi \geqslant 2$ case from Equation (1) in the calculation of Equation (19). Analyzing the three cases where the limits of integration in (19) satisfy $(\alpha/r_s)(r_c-r_s) \leqslant 0$, $0 < (\alpha/r_s)(r_c-r_s) < \pi/M$, and $(\alpha/r_s)(r_c-r_s) \geqslant \pi/M$, respectively, yields

$$d_b = \begin{cases} \tilde{r}_c - \dfrac{M}{\pi}(\alpha\tilde{r}_c - \tilde{r}_s)\ln\left(1+\dfrac{\pi}{(\alpha-1)M}\right) & \text{if } \tilde{r}_c \leqslant \tilde{r}_s, \\ \tilde{r}_c + \dfrac{\alpha M(\tilde{r}_c-\tilde{r}_s)^2(\alpha+2)}{2\pi\tilde{r}_s(\alpha+1)} - \dfrac{\alpha M\tilde{r}_c(\tilde{r}_c-\tilde{r}_s)}{\pi\tilde{r}_s} \\ \quad - \dfrac{M}{\pi}(\alpha\tilde{r}_c-\tilde{r}_s)\ln\left(\dfrac{\alpha-1+\pi/M}{\alpha-1+(\alpha/\tilde{r}_s)(\tilde{r}_c-\tilde{r}_s)}\right) & \text{if } \tilde{r}_s < \tilde{r}_c < \tilde{r}_s(1+\pi/(\alpha M)), \\ \dfrac{\tilde{r}_c - \tilde{r}_s(1-\pi/2M)}{\alpha+1} & \text{if } \tilde{r}_c \geqslant \tilde{r}_s(1+\pi/(\alpha M)), \end{cases} \tag{20}$$

and we set $t_b = d_b/R$ in computing the service rate $\mu$ of the $M/M/1/2$ queue. A nonreneging customer who travels $d_b$ after the chase begins is caught at radial location $\tilde{r}_c - d_b$.

Taken together, the expected damage is approximated by

$$E[D] = qp_0[b - aR(1-t_e)] + p_1\left(\frac{r_1 b}{r_1+\mu} + \frac{\mu q[b - a(\tilde{r}_c - d_b)]}{r_1+\mu}\right) + p_2 b. \quad (21)$$

## 4. Computational Study

We set $R = 50$ miles, $\alpha = 1.5$, $m_s = 0.5$ hr, $\sigma_s = 0.05$ hr, and $q = 0.9$. As noted in §2, we also set $b = 10$ and $a = 9/R$ to maintain a 1-to-10 damage scale. First, we assess the accuracy of the approximation in §3 using a simulation model, which is described in detail in §B of the online companion. For different numbers of servers ($M$), Figure 1 plots the simulated mean damage from the model in §2 and the approximate mean damage in Equation (21) versus the arrival rate $\lambda$. For the case of $M = 10$ servers, we also plot the traffic intensity $\rho$, which we define by

$$\rho = \frac{\lambda}{M}(m_s + m_t), \quad (22)$$

where $m_s$ is the mean on-site service time and $m_t$ is the analytical approximation to the mean travel component of service given by Equation (11). Simulation results (see §B of the online companion) reveal that Equation (22) is an accurate approximation to the actual utilization rate of the model in §2. Figure 2 in the online companion displays the probability that a terrorist reaches the target versus the arrival rate; if we denote this probability by $p_T$, then the probability that a terrorist successfully detonates is $p_T + (1-p_T)q$.

**Figure 1.** For various numbers of interdiction vehicles, the simulated (see §B of the online companion and §2, with bars denoting the 95% confidence intervals) and analytical (Equation (21)) mean damage vs. $\lambda$ and vs. $\rho$.

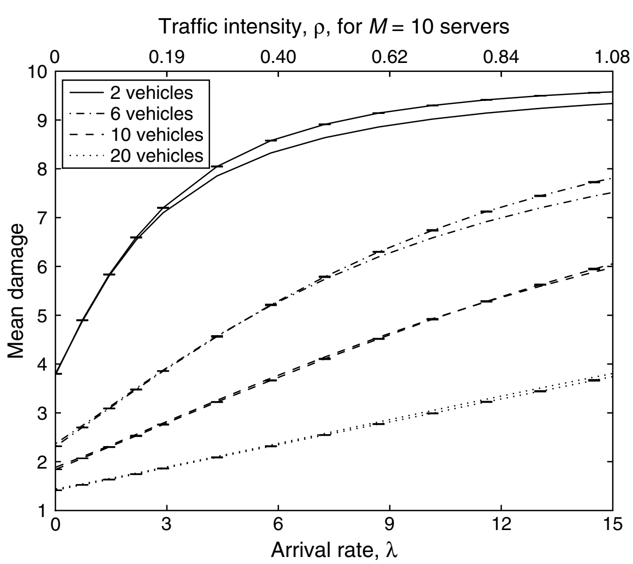


The simulated values were derived by averaging over 10 simulations, each with 100,000 arrivals and truncating the statistics from the first 10,000. Equation (21) accurately approximates the simulated values for all traffic intensities. A detailed investigation (data not shown) reveals that for the particular set of parameter values under consideration, the mean chase time in (5) (even though our optimal resting location underestimates the true value by approximately $0.05R$), the steady-state probabilities in (6), the reneging probability in (14), the radial location in (18), and the mean distance traveled in (20) are all highly accurate for $\rho \leqslant 0.3$. For $\rho > 0.3$, the approximation overestimates the impact of customers who reach the center and underestimates the impact of customers who are served after entering a busy queue. These two effects offset each other, yielding the highly accurate damage estimate shown in Figure 1. The reason for these misestimations in heavier traffic is primarily because the assumption above Equation (12)—if a customer arrives to a queue with one customer in service, then the customer currently in service arrived to an empty queue—is no longer valid. This causes underestimation of $T_1$ in Equation (12), overestimation of $P_r$ in Equation (14), and underestimation of $d_b$ in (20), and the latter two directly impact the steady-state probabilities in Equation (6). Figure 2 in the online companion illustrates the accuracy of the estimate of the number of customers that reach the center.

Figure 1 suggests that many servers are needed to maintain moderate damage when the arrival rate exceeds 3/hr. This plot also suggests that the mean damage is approximately linear in the traffic intensity (for $\rho \leqslant 1$) for a fixed number of servers. In very light traffic, we can use (8) to approximate the mean damage in (21) by the simpler expression

$$E[D] = q\left(b - aR\left(1 - \frac{(\alpha+2)\pi}{2(\alpha+1)(\alpha M + \pi)}\right)\right) \quad (23)$$

$$= 0.9 + \frac{89}{7.5M + 15.7}, \quad (24)$$

with the parameter values used in our computational study. Figure 1 shows that Equation (24) accurately predicts the $y$-intercept of these curves. Inverting Equation (23) provides a heuristic (but loose if $\rho > 0.1$) lower bound for the number of servers required to maintain the mean damage to the level $E[D]$:

$$M \geqslant \frac{qaR(\alpha+2)\pi}{2\alpha(\alpha+1)[E[D] - q(b - aR)]} - \frac{\pi}{\alpha}. \quad (25)$$

Finally, we simulate a system in which servers can aid each other. More specifically, we alter the simulation model so that if the server in a wedge is busy at the time of a customer arrival to his wedge, then this customer is served by the closest idle server. If all servers are busy when the customer arrives, then the first available server that can catch the customer before he reaches the target will serve the customer. If no servers can successfully catch the customer, the customer proceeds directly to the target. We assume that when the server is idle, he returns to his optimal resting location (as in our primary model).

Figure 3 in the online companion shows that the pooled system can reduce the mean damage by nearly a third (e.g., from three to two) in light traffic. Even though a smaller fraction of customers reach the center in the pooled system, the unpooled system outperforms it in heavy traffic because the customers that are caught in the pooled model are stopped much closer to the center than the corresponding customers that are caught in the unpooled queue, and because servers in the pooled system spend too much time traveling, and hence are out of position for customers arriving to their wedge. There are many variations on such a pooled system (e.g., subset of servers that help each other, non-FIFO discipline, dynamic server resting locations), but we do not pursue them here.

## 5. Concluding Remarks

Our analysis suggests that rapid interdiction requires many interdiction vehicles if the arrival rate of alarm-generating vehicles to the city is greater than 3/hr, and Equations (21) and (25) provide an accurate estimate and a back-of-the-envelope lower bound, respectively, for the number of interdiction vehicles required to maintain the mean damage below any specified level. However, readers should bear in mind that this spatial interdiction model is just a caricature of an actual highway system. Although our model is appropriate for its intended use as a rough-cut feasibility study, more refined recommendations would require a specific city's network to be modeled and would need to incorporate other operational issues, such as cooperation among vehicles (see §4 for a start in this direction) and nonhomogeneous arrival rates.

A cursory understanding of radiation detection is required to put into perspective the model's key parameter, which is the arrival rate of alarm-generating vehicles. Radiation sensors measure neutrons and gamma rays. Uranium and plutonium, which are the two possible sources for a nuclear weapon, emit both neutrons and gamma rays, and are among the only substances that emit neutrons. Dirty bombs, along with many other legal items (e.g., kitty litter, ceramic tiles, bananas) emit gamma rays. Currently deployed technologies aggregate the gamma emissions rather than look at the emissions along the entire energy spectrum, leading to a high false-positive rate (Rooney 2005 estimates it at 40%). In contrast, the vehicle arrival rate into a large city can be 100,000/hr, and a mean damage of 3 (on a 1-to-10 scale) can be maintained by $M = 20$ interdiction vehicles only if the false-positive probability is less than $10^{-4}$ (this corresponds to $\lambda = 10$ in Figure 1). Hence, current gamma-ray detection technology is impractical for this application, precluding the detection of dirty bombs; however, spectroscopic gamma-ray detectors, which may be capable of reducing the false-positive probability, are just starting to be deployed at U.S. ports (Lipton 2006). Neutron detectors appear capable of detecting plutonium, but not uranium (Huizenga 2005). Taken together, a detection-interdiction system using current technology appears capable of detecting a plutonium bomb, but not a uranium bomb or a dirty bomb, although a thick wall of sensors (i.e., a vehicle would have to pass through many sensors) might generate a slight increase in the overall detection probability of a uranium weapon.

In a companion paper (Wein and Atkinson 2007), we embed three models into a Stackelberg game: a sensor model first developed in Wein et al. (2006), which determines the detection probability and the false-positive probability as a function of the neutron threshold level of the sensor; an optimal stopping problem for the terrorist (whether to proceed directly to the circle center or to detonate at any point along his route, based on a Bayesian update of the detection probability of sensors carried out in Atkinson et al. 2008); and the spatial interdiction model analyzed here. In this game, the U.S. government (as the leader) chooses the neutron threshold level, the thickness of the wall sensors (i.e., how many sensors the terrorist needs to pass through), and the number of interdiction vehicles to minimize the expected damage inflicted by a terrorist, subject to a budget constraint on the annual cost of sensors and interdiction vehicles. The terrorist (as the follower) observes the wall thickness and solves the optimal stopping problem with the goal of maximizing the expected damage. Although using a much different model in a different setting—pedestrian suicide bombers—Kaplan and Kress (2005) analyze a model that also takes into account sensors, terrorist behavior, and interdiction.

## 6. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.pubs.informs.org/ecompanion.html.

## Acknowledgments

The second author thanks Tom Edmunds and Richard Wheeler for helpful conversations. The authors thank the reviewers for their valuable comments. This research was primarily supported by Lawrence Livermore National Laboratory, Project #B529238. The first and second authors were also supported, respectively, by an Abbott Laboratories Stanford Graduate Fellowship and the Center for Social Innovation, and the Graduate School of Business, Stanford University.

## References

Atkinson, M. P., Z. Cao, L. M. Wein. 2008. Optimal stopping analysis of a radiation detection system to protect cities from a nuclear terrorist attack. *Risk Anal.* Forthcoming.

Bertsimas, D. J., G. van Ryzin. 1993. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Oper. Res.* **41** 60–76.

Bunn, M., A. Wier, J. P. Holdren. 2003. Controlling nuclear warheads and materials: A report card and action plan. Project on Managing the Atom, John F. Kennedy School of Government, Harvard University, Cambridge, MA.

Flynn, S. E. 2000. Beyond border control. *Foreign Affairs* **79** 57–65.

Glasstone, S., P. J. Dolan. 1977. *The Effects of Nuclear Weapons*, 3rd ed. U.S. Dept. of Defense, Energy Research and Development Administration, Washington, D.C.

Gross, D., C. M. Harris. 1985. *Fundamentals of Queueing Theory*. John Wiley & Sons, New York.

Heyman, D. P., M. J. Sobel. 1982. *Stochastic Models in Operations Research*, Vol. 1. McGraw Hill, New York.

Huizenga, D. 2005. Detecting nuclear weapons and radiological materials: How effective is available technology? Testimony Before the Subcommittee on Prevention of Nuclear and Biological Attacks and the Subcommittee on Emergency Preparedness, Science and Technology, The House Committee on Homeland Security, June 21, Washington, D.C.

Kaplan, E. H., M. Kress. 2005. Operational effectiveness of suicide-bomber-detector schemes: A best-case analysis. *Proc. Natl. Acad. Sci. USA* **102** 10399–10404.

Larson, R. C., A. R. Odoni. 1981. *Urban Operations Research*. Prentice-Hall, Englewood Cliffs, NJ.

Lipton, E. 2006. U.S. to spend $1.2 billion on detecting radiation. *The New York Times* (July 15) A10.

Lipton, E. 2007. New York to test ways to prevent nuclear terror. *The New York Times* (Feburary 9) A1.

Pan, F., W. S. Charlton, D. P. Morton. 2003. A stochastic program for interdicting smuggled nuclear material. D. L. Woodruff, ed. *Network Interdiction and Stochastic Programming*, Chapter 1. Kluwer, Boston, 1–19.

Rooney, B. 2005. Detecting nuclear weapons and radiological materials: How effective is available technology? Testimony Before the Subcommittee on Prevention of Nuclear and Biological Attacks and the Subcommittee on Emergency Preparedness, Science and Technology, The House Committee on Homeland Security, June 21, Washington, D.C.

Stana, R. M. 2004. Summary of challenges faced in targeting oceangoing cargo containers for inspection. U.S. General Accounting Office Report GAO-04-557T, March 31, Washington, D.C.

U.S. Department of Homeland Security. 2005. Fact sheet: Domestic nuclear detection office. Retrieved September 23 http://www.dhs.gov/dhspublic/display?theme=43&content=4474&print=true.

Washburn, A., K. Wood. 1995. Two-person zero-sum games for network interdiction. *Oper. Res.* **43** 243–251.

Wein, L. M., M. P. Atkinson. 2007. The last line of defense: Designing and managing radiation sensor arrays around cities. *IEEE Trans. Nuclear Sci.* **54** 654–669.

Wein, L. M., A. H. Wilkins, M. Baveja, S. E. Flynn. 2006. Preventing the importation of illicit nuclear materials in shipping containers. *Risk Anal.* **26** 1377–1393.

Wollmer, R. D. 1964. Removing arcs from a network. *J. Oper. Res. Soc. America* **12** 934–940.

Woodruff, D. L. 2003. *Network Interdiction and Stochastic Programming*. Kluwer, Boston.

OPERATIONS RESEARCH
DOI 10.1287/opre.1070.0490ec
pp. ec1–ec1



**e-companion**

**ONLY AVAILABLE IN ELECTRONIC FORM**

# Electronic Companion—"Spatial Queueing Analysis of an Interdiction System to Protect Cities from a Nuclear Terrorist Attack" by Michael P. Atkinson and Lawrence M. Wein, *Operations Research* 2008, DOI 10.1287/opre.1070.0490.

# Online Companion

Section A contains the proofs of Proposition 1. The simulation study is described in §B. Figures 2-3 are also in this Online Companion.

## A Proof of Proposition 1

We first look at the $\psi \geq 2$ case. In this scenario the server will never travel along arcs. To travel from one location $(r_s, \theta_s)$, to another location $(r_s, \theta_1)$, with the same radial component, the server has two options. The server can travel a distance $r_s|\theta_s - \theta_1| = r_s\psi$ along an arc, or the server can travel a distance $2r_s$ by first traveling from his current location, $(r_s, \theta_s)$, to the center, and then traveling from the center to his desired location $(r_s, \theta_1)$. Thus if $\psi \geq 2$ it is faster to travel along rays into and out of the center of the circle rather than along arcs.

In this case, when the server chases a customer starting from $(r_c, \theta_c)$ he first travels a distance $r_s$ to the center, and after traveling a distance $\frac{r_s}{\alpha}$, the customer is located at $(r_c - \frac{r_s}{\alpha}, \theta_c)$. Starting from the center the server drives along the same ray as the customer, who travels an additional $\frac{1}{\alpha+1}\left(r_c - \frac{r_s}{\alpha}\right)$ before being caught. Therefore the total distance traveled by the customer is $\frac{r_c+r_s}{\alpha+1}$.

We now assume $\psi < 2$ and first analyze the arc/ray approach. While moving along the arc from $\theta_s$ to $\theta_c$, the server travels the distance $r_s\psi$ and the customer travels $\frac{r_s\psi}{\alpha}$. If $r_c - \frac{r_s\psi}{\alpha} \geq r_s$ the customer is traveling toward the server and he travels an additional $\frac{1}{\alpha+1}\left(r_c - \frac{r_s\psi}{\alpha} - r_s\right)$ distance before being caught.

If $r_c - \frac{r_s\psi}{\alpha} < r_s$ then the customer is driving away from the server. If the customer travels an additional $d'$ before interdiction the server travels $d' + r_s - r_c + \frac{r_s\psi}{\alpha}$, yielding the equation $\alpha d' = d' + r_s - r_c + \frac{r_s\psi}{\alpha}$, or $d' = \frac{1}{\alpha-1}\left(r_s - r_c + \frac{r_s\psi}{\alpha}\right)$. Hence, the total distance the customer travels under the arc/ray strategy is

$$d(r_c, r_s, \psi) = \begin{cases} \frac{r_c - r_s + r_s\psi}{\alpha+1} & \text{if} \quad r_c - \frac{r_s\psi}{\alpha} \geq r_s; \\ \frac{r_s - r_c + r_s\psi}{\alpha-1} & \text{if} \quad r_c - \frac{r_s\psi}{\alpha} < r_s. \end{cases} \tag{26}$$

Turning to the ray/arc approach, if the server interdicts the customer at a radius greater than $r_s$, then the server first moves toward the circumference a distance of $d'$ and then moves $(r_s + d')\psi$ along the arc to catch the customer. Meanwhile, the customer travels a distance of $\frac{d' + (r_s + d')\psi}{\alpha}$. Equating their respective radius locations gives $r_c - \frac{d' + (r_s + d')\psi}{\alpha} = r_s + d'$, yielding $d' = \frac{\alpha r_c + r_s}{\alpha+1+\psi} - r_s$. The customer travels a distance $r_c - (r_s + d')$, which equals $\frac{r_c - r_s + r_c\psi}{\alpha+1+\psi}$.

Similarly, if the server catches the customer at a radius less than $r_s$, the server first moves toward the circle center a distance of $d'$ and then moves $(r_s - d')\psi$ along the circular arc to reach the customer. This suspect travels a distance of $\frac{d' + (r_s - d')\psi}{\alpha}$. The unknown distance $d'$ satisfies $r_c - \frac{d' + (r_s - d')\psi}{\alpha} = r_s - d'$, and hence $d' = r_s - \frac{\alpha r_c - r_s}{\alpha-1+\psi}$, and the customer travels a distance of $r_c - (r_s - d')$, or $\frac{r_s - r_c + r_c\psi}{\alpha-1+\psi}$. So the total distance the customer travels under the ray/arc strategy is

$$d(r_c, r_s, \psi) = \begin{cases} \frac{r_c - r_s + r_c\psi}{\alpha+1+\psi} & \text{if} \quad r_c - \frac{r_s\psi}{\alpha} \geq r_s; \\ \frac{r_s - r_c + r_c\psi}{\alpha-1+\psi} & \text{if} \quad r_c - \frac{r_s\psi}{\alpha} < r_s. \end{cases} \tag{27}$$

We now need to compare equations (26)-(27) to determine which is better. If the radius $(r_c - \frac{r_s\psi}{\alpha})$ is greater than $r_s$, then assuming $\psi \neq 0$ (if $\psi = 0$ the methods are equivalent)

$$\begin{aligned} & r_s\psi \leq \alpha(r_c - r_s) \\ \iff \quad & r_c - r_s + (\alpha + 1 + \psi)r_s \leq (\alpha + 1)r_c \\ \iff \quad & \psi(r_c - r_s) + (\alpha + 1 + \psi)r_s\psi \leq (\alpha + 1)r_c\psi \\ \iff \quad & (\alpha + 1 + \psi)(r_c - r_s + r_s\psi) \leq (\alpha + 1)(r_c - r_s + r_c\psi) \\ \iff \quad & \frac{r_c - r_s + r_s\psi}{\alpha + 1} \leq \frac{r_c - r_s + r_c\psi}{\alpha + 1 + \psi}, \end{aligned}$$

and the arc/ray strategy is better. Similarly, if $(r_c - \frac{r_s\psi}{\alpha})$ is less than $r_s$ then

$$\begin{aligned} & r_s\psi > \alpha(r_c - r_s) \\ \Longleftrightarrow\quad & r_s - r_c + (\alpha - 1 + \psi)r_s > (\alpha - 1)r_c \\ \Longleftrightarrow\quad & \psi(r_s - r_c) + (\alpha - 1 + \psi)r_s\psi > (\alpha - 1)r_c\psi \\ \Longleftrightarrow\quad & (\alpha - 1 + \psi)\,(r_s - r_c + r_s\psi) > (\alpha - 1)(r_s - r_c + r_c\psi) \\ \Longleftrightarrow\quad & \frac{r_s - r_c + r_s\psi}{\alpha - 1} > \frac{r_s - r_c + r_c\psi}{\alpha - 1 + \psi}, \end{aligned}$$

and the ray/arc strategy prevails.

Finally, we will show that if the interdiction vehicle takes a different strategy than ray/arc or arc/ray (e.g., ray/arc/ray/arc/ray where the vehicle makes several turns and travels on multiple different arcs) then the customer will not travel a shorter distance into the system. This analysis only applies for $\psi < 2$, because we have already shown that for $\psi \geq 2$ it is suboptimal to travel along arcs, and thus since the server has only one possible interdiction path that does not involve traveling on any arcs it must be the optimal path. First let us assume the customer is caught in front of the server (i.e, $r_c - \frac{r_s\psi}{\alpha} \geq r_s$). Assuming the arc/ray strategy is used, the interdiction point occurs at $(r^*, \theta_c)$, where $r^* = \frac{\alpha r_c + r_s - r_s\psi}{\alpha+1}$. The total distance traveled by the interdiction vehicle is $r^* - r_s + r_s\psi$. Now let us assume a different path consisting of several ray and arc segments. The interdiction point for this path would occur at some location $(\hat{r}, \theta_c)$. We want to show that $r^*$ is the maximum interdiction radius possible, so we will assume $\hat{r} > r^*$ and derive a contradiction. We further assume that this path makes $n$ turns at radii $r_i$, $i = 1, \ldots, n$ and traverses the arc on $r_i$ for $\psi_i$ radians, where $\sum_{i=1}^{n} \psi_i = \psi$ and $r_i \in [r_s, \hat{r}]$. Thus the interdiction vehicle travels a distance of $\hat{r} - r_s + \sum_{i=1}^{n} \psi_i r_i$. However if $\hat{r} > r^*$ and $\sum_{i=1}^{n} \psi_i r_i \geq \sum_{i=1}^{n} \psi_i r_s = r_s\psi$, then the interdiction vehicle must travel further than he did to catch the suspect at $r^*$ using the arc/ray strategy, which implies the customer also travels a greater distance, and thus the interdiction point is less than $r^*$, which yields the contradiction.

Next let us assume the customer is caught behind the server (i.e, $r_c - \frac{r_s \psi}{\alpha} < r_s$). Assuming the ray/arc strategy is used, the interdiction point occurs at $(r^*, \theta_c)$, where $r^* = \frac{\alpha r_c - r_s}{\alpha - 1 + \psi}$. The total distance traveled by the interdiction vehicle is $r_s - r^* + r^* \psi$. Now let us assume a different path consisting of several ray and arc segments. The interdiction point for this path would occur at some location $(\hat{r}, \theta_c)$. We want to show that $r^*$ is the maximum interdiction radius possible, so we will assume $\hat{r} > r^*$ and derive a contradiction. We further assume that this path makes $n$ turns at radii $r_i$, $i = 1, \ldots, n$ and traverses the arc on $r_i$ for $\psi_i$ radians, where $\sum_{i=1}^{n} \psi_i = \psi$ and $r_i \in [\hat{r}, r_s]$. Thus the interdiction vehicle travels a distance of $r_s - \hat{r} + \sum_{i=1}^{n} \psi_i r_i$. However, because $r_i \geq \hat{r}$, it is suboptimal for $n > 1$ (take more than one turn), and $r_i > \hat{r}$ since we can achieve the distance $r_s - \hat{r} + \hat{r}\psi$. This is just the distance traveled from $r_s$ to $\hat{r}$ using the ray/arc strategy, and so it is impossible for $\hat{r} > r^*$ because $r^*$ is the interdiction point for the ray/arc strategy and as the analysis in the paragraph above (27) shows, the ray/arc approach has a unique interdiction point. Thus we have arrived at a contradiction and have completed the proof of Proposition 1.

# B The Simulation Model

In the computational study, we analyze only one slice. First we simulate 100,000 arrivals by generating exponential random variables with mean $\frac{M}{\lambda}$. We denote the interarrival time of customer $i$ as $\tau_i$ and the arrival time of customer $i$ as $t_i$, so that $t_n = \sum_{i=1}^{n} \tau_i$. We assume that these customers arrive on the edge of the slice (at distance $R$ from the center) uniformly between $\theta_i \in [0, \frac{2\pi}{M}]$. At the beginning of the simulation, the interdiction vehicle is located at its optimal resting location $(r^*, \theta^*)$ (described in §2). The radial coordinate, $r^*$, of the optimal resting location is given by equation (2) in the main text (it is exact, not the approximation in (3)) and it is determined numerically, and the angular coordinate, $\theta^*$, is $\frac{\pi}{M}$.

In the first step of the simulation, the interdiction vehicle serves the first customer arriving to the queue at location $(R, \theta_1)$ at time $t_1$. As soon as the customer arrives to the system, he starts driving straight toward the target located at the center of the circle, and the interdiction vehicle immediately sets off from his optimal resting location to serve this customer.

The time and distance traveled by this customer before interdiction, which we denote as $d_1$, is given by equation (1) in the main text (where $r_s = r^*$, $r_c = R$ and $\psi = |\theta_1 - \frac{\pi}{M}|$). The customer receives on-site service time of $o_1$, which we determine by generating a normal random variable with mean 0.5 hr and standard deviation 0.05 hr. Thus the total service time for the first customer is $s_1 = d_1 + o_1$ and the customer is stopped a distance $r_1 = R - d_1$ from the target. After service is complete, we assume with probability $q$ the customer detonates a bomb causing damage $D_1 = b - ar_1$, and with probability $1 - q$ the plot was successfully foiled causing damage 0. Service for the first customer is considered finished at time $t_1 + s_1$, at which point the server is free to serve other customers in the system.

We next look at the second customer who arrives to the system at time $t_2$ at location $(R, \theta_2)$. There are several different scenarios involving this customer. If $t_2 < t_1 + s_1$ then the second customer arrives to the system before the first customer completes service. Furthermore, if $t_2 + R < t_1 + s_1$ then the second customer reaches the origin before the first customer completes service (since customers start traveling toward the target as soon as they arrive to the system). In this case, the second customer causes damage $D_2 = b$ because there is no interdiction.

If $t_2 < t_1 + s_1$ but $t_2 + R \geq t_1 + s_1$, then the second customer is in the system traveling toward the target when service on the first vehicle has completed, but the second customer has not yet reached the origin, and is located at $R - (t_1 + s_1 - t_2)$ from the origin. In this scenario one of two things can happen: either the server can still catch this customer before he reaches the center or he cannot. After serving the first customer, the server is located at $(r_1, \theta_1)$. If $R - (t_1 + s_1 - t_2) < \frac{r_1}{\alpha}$, the server cannot catch the second customer before he reaches the center and thus makes no attempt to catch him, and the customer causes damage $D_2 = b$. If $R-(t_1+s_1-t_2) \geq \frac{r_1}{\alpha}$ the server can catch the customer, and the customer travels a distance $d_2$ during the pursuit, given by equation (1) in the main text (where $r_s = r_1$, $r_c = R - (t_1 + s_1 - t_2)$ and $\psi = |\theta_1 - \theta_2|$). The second customer is stopped a distance $r_2 = R - (t_1 + s_1 - t_2) - d_2$ from the origin receiving a random normal service time $o_2$, for a total service time of $s_2 = d_2 + o_2$. After service is complete, we assume with probability $q$ the customer detonates a bomb causing damage $D_2 = b - ar_2$, and with probability $1-q$ the plot was successfully foiled causing damage 0. Service for the second customer is considered finished at time $t_1 + s_1 + s_2$, at which point the server is free to serve other customers in the system.

If $t_2 \geq t_1 + s_1$ then the second customer arrives to an empty queue, and again we have different cases to analyze. After finishing serving the first customer at time $t_1 + s_1$ at location $(r_1, \theta_1)$ the server starts back to his optimal resting location at $(r^*, \frac{\pi}{M})$. The server takes the fastest route back to the optimal resting location (under the constraint that he travels along rays and arcs) and it takes him $t_{opt,1} = \frac{|r^*-r_1|+|\theta_1-\frac{\pi}{M}|\times\min(r^*,r_1)}{\alpha}$ time units to get back there (since we analyze $M \geq 2$, we have $|\theta_1 - \frac{\pi}{M}| \leq 2$ and thus the interdiction vehicle will travel along arcs and we do not have to consider the first case in equation (1)). If $t_2 \geq t_1+s_1+t_{opt,1}$ then the server gets back to his optimal resting location before the second customer arrives to the system. In this case, analyzing the behavior of the second customer is equivalent to the first customer because the second customer arrives to an empty system with the server in the optimal resting location. If, however, $t_2 < t_1 + s_1 + t_{opt,1}$, then the server is located at some generic location in the system $(r_s, \theta_s)$ along the path back to the optimal resting location when the second customer arrives. The server immediately leaves from this location, $(r_s, \theta_s)$, and chases down the customer. The analysis in this case is the same as for the first customer, the only difference is that rather than starting the chase from his optimal location $(r^*, \frac{\pi}{M})$, the server leaves from the arbitrary location $(r_s, \theta_s)$.

The rest of the customers are analyzed in a similar way. If a customer cannot be caught and reaches the center before being served then the server immediately determines the next customer in the queue who can be served and acts to optimally serve that customer.

So to summarize, after successfully servicing some customer, say customer $i$, at time $\tilde{t}_i$ (the ending service time of an arbitrary customer can be complicated) the server is located at $(r_i, \theta_i)$ and there are five possible scenarios for customer $i + 1$:

1. $t_{i+1} + R < \tilde{t}_i$. Customer $i+1$ reaches the center before service is complete on customer $i$. In this case customer $i + 1$ receives no service. The server determines if customer $i + 2$ can receive service, and if so he goes about serving that customer. If not, the server determines the next customer he can serve.

2. $t_{i+1} < \tilde{t}_i \leq t_{i+1} + R,\ R - (\tilde{t}_i - t_{i+1}) < \frac{r_i}{\alpha}$. Customer $i + 1$ enters a busy system, the server is free before the customer reaches the center, but the server still cannot catch the customer before the customer reaches the target. In this case customer $i + 1$ receives no service. The server determines if customer $i+2$ can receive service, and if so he goes about serving that customer. If not, the server determines the next customer he can serve.

3. $t_{i+1} < \tilde{t}_i \leq t_{i+1} + R,\ R - (\tilde{t}_i - t_{i+1}) \geq \frac{r_i}{\alpha}$. Customer $i + 1$ enters a busy system, the server is free before the customer reaches the center, and the server can catch the customer before the customer reaches the target. In this case the server immediately leaves from $(r_i, \theta_i)$, catches and serves customer $i + 1$.

4. $\tilde{t}_i \leq t_{i+1}$, $t_{i+1} \geq \tilde{t}_i + t_{opt,i}$. Customer $i+1$ enters an empty system, and the server has arrived back to his optimal resting location by the time customer $i+1$ arrives to the system. In this case the server leaves from $(r^*, \frac{\pi}{M})$, catches and serves customer $i+1$.

5. $\tilde{t}_i \leq t_{i+1}$, $t_{i+1} < \tilde{t}_i + t_{opt,i}$. Customer $i+1$ enters an empty system, and arrives before the server reaches his optimal resting location. The server immediately leaves from his general location $(r_s, \theta_s)$, catches and serves customer $i+1$.

After looping through all 100,000 customers, we have how much damage, $D_i$, each caused, as well as many other values related to the customer (e.g., whether the customer was served, the service location of the customer, the total service time of the customer, whether a customer entered a busy or empty queue, etc.). Since we are interested in the steady-state behavior of the system, we drop the first 10% and only look at the last 90,000 arrivals. We take the average of the $D_i$ from these to compute an estimate of the expected damage from the system. We also compute the average for other values of interest (such as distance traveled) to compare with the analytic approximations from our M/M/1/2 model. Finally we repeat the simulation 10 times. The average damage over the 10 simulations is plotted in Figure 1. The error bars in Figure 1 are computed using the sample variance of the 10 simulation runs.

Finally, in making Figures 1, 2 and 3, the top horizontal axis is the analytical $\rho$ given by equation (22) in the main text. The actual value of $\rho$ for a simulation run is given by $\rho = \frac{\lambda}{M} \frac{1}{N} \sum_{i=1}^{N} s_i$ where $N$ is the number of customers we are simulating. Hence, in Figures 1, 2 and 3, the true value of $\rho$ in the simulation may be slightly different than the value given on the top horizontal axis, but in our experience the difference between the simulated and analytical values is less than 1% for $\rho < 0.3$ and less than 5% for larger values of $\rho$. We have no analytic expression for $\rho$ in the pooled queue model used for Figure 3, but we use equation (22) in the main text to derive the top horizontal axis for comparison purposes with the unpooled model.

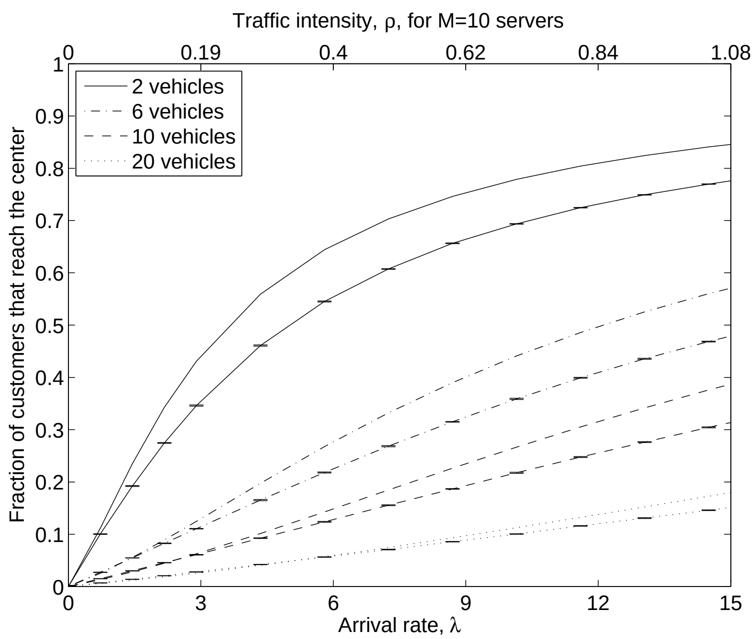


Figure 2: For various numbers of interdiction vehicles, the simulated (§B of the Online Companion and §2, with bars denoting the 95% confidence intervals) and analytical ($p_2+\frac{p_1 r_1}{r_1+\mu}$ from §3) fraction of customers that reach the center vs. $\lambda$ and vs. $\rho$.

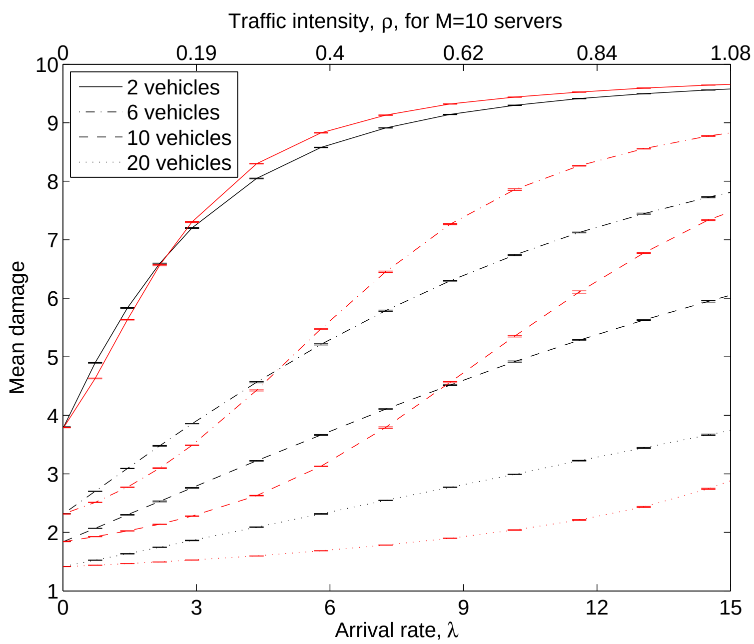


Figure 3: For various numbers of interdiction vehicles, the simulated mean damage vs. $\lambda$ and vs. $\rho$ for the unpooled system (i.e., as in Figure 1 in the main text) and the pooled system (red) in which vehicles can travel to other wedges.