

The Last Line of Defense: Designing Radiation Detection-Interdiction Systems to Protect Cities From a Nuclear Terrorist Attack

Lawrence M. Wein and Michael P. Atkinson

Abstract—We formulate and solve an optimization problem in which a terrorist is attempting to drive a nuclear weapon toward a city center, but needs to travel through an array of imperfect neutron radiation sensors that form a wall around the periphery of the city. A fleet of interdiction vehicles are available to chase, and attempt to interdict, vehicles that set off a sensor alarm. In our model, the government chooses the thickness (in terms of number of sensors) of the radiation wall, the neutron threshold in the sensors, and the number of interdiction vehicles to minimize the expected damage inflicted by a terrorist, subject to a budget constraint on sensors and interdiction vehicles. The terrorist observes the wall thickness and at each node he updates his likelihood of passing through a sensor without triggering an alarm and decides whether to proceed through the sensor or stop and detonate the bomb. Our results suggest that for an annual cost ranging from several million dollars to several tens of millions of dollars, depending upon the city's roadway topology, a single layer of sensors placed tens of miles from the city center and 10–20 dedicated interdiction vehicles could mitigate the damage from an unshielded or lightly-shielded plutonium weapon, but not from a uranium weapon or a radiological dispersal device.

Index Terms—Game theory, optimal stopping, queueing theory, radiation detection.

I. INTRODUCTION

IN the aftermath of the September 11, 2001 attacks on the World Trade Center and the Pentagon, the U.S. Government's most feared scenario is that terrorists acquire and detonate a nuclear weapon in a major U.S. city. A Harvard research group commissioned by the Nuclear Threat Initiative has deemed this threat, which they estimate could cause a half-million deaths and one trillion dollars in direct economic damage, to be "real and urgent" [1]. Efforts by the Cooperative Threat Reduction Program to secure the world's stockpiles of nuclear weapons and materials, particularly in the former Soviet Union, have moved slowly, and the majority of nuclear material remains vulnerable to theft [1]. Well-organized terrorist groups are capable of making at least a crude nuclear weapon [1], and should

have little trouble smuggling nuclear material or weapons into a U.S. port in a shipping container [2], [3].

This paper formulates and analyzes a mathematical model for our last line of defense, which is to detect an assembled weapon as it is driven into a U.S. city and to provide timely and effective interdiction. Plans are under way to develop such a detection-interdiction system under the auspices of the newly-created Domestic Nuclear Detection Office [4], and our goal here is to perform a rough-cut feasibility analysis.

II. THE MODEL

A. Model Overview

The model uses three methodologies that may be unfamiliar to journal readers: game theory (in particular, a Stackelberg game), stochastic dynamic programming (in particular, an optimal stopping problem), and queueing theory. These methodologies will be briefly discussed as we give a broad overview of the model.

Game theory is an optimization framework with more than one (and often two) decision makers, typically with opposing objectives [5]. In many games, the two players make their decisions simultaneously, such as in the simplest of all games, the Prisoner's Dilemma. However, in many practical situations, it is more realistic to assume that the players move sequentially, such games are referred to as Stackelberg games, where the "leader" makes his decisions and then the "follower", after observing some of the leader's actions, makes his decisions. Stackelberg games are particularly appropriate for homeland security problems, where it is typically assumed (for reasons of realism and conservatism) that the government is the leader and the terrorist is the follower [6]–[8]. We follow this approach here: the government chooses the thickness of the radiation sensor wall (i.e., the number of sensors a terrorist would need to pass through before reaching his target), the number of interdiction vehicles, and the neutron threshold level that triggers each sensor, with the goal of minimizing the expected damage from a detonated (plutonium or uranium) bomb. The terrorist then decides how far to proceed through the network before detonating the bomb in order to maximize the expected damage. If the terrorist triggers a sensor then an interdiction vehicle (if it is not busy chasing and investigating non-terrorist vehicles that set off a sensor alarm) chases the terrorist's vehicle, at which point interdiction is successful (i.e., the bomb is not detonated) with a specified probability.

We describe our model in terms of three interconnected sub-models (Fig. 1) and all parameters and their base-case values are

Manuscript received January 20, 2006; revised January 3, 2007. This work was supported by the Lawrence Livermore National Laboratory, Project B529238. L. M. Wein was supported by the Center for Social Innovation, Graduate School of Business, Stanford University. M. P. Atkinson was supported by an Abbott Laboratories Stanford Graduate Fellowship.

L. M. Wein is with the Graduate School of Business, Stanford University, Stanford, CA 94305 USA; (e-mail: lwein@stanford.edu)

M. P. Atkinson is with the Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: mpa33@stanford.edu)

Digital Object Identifier 10.1109/TNS.2007.897829

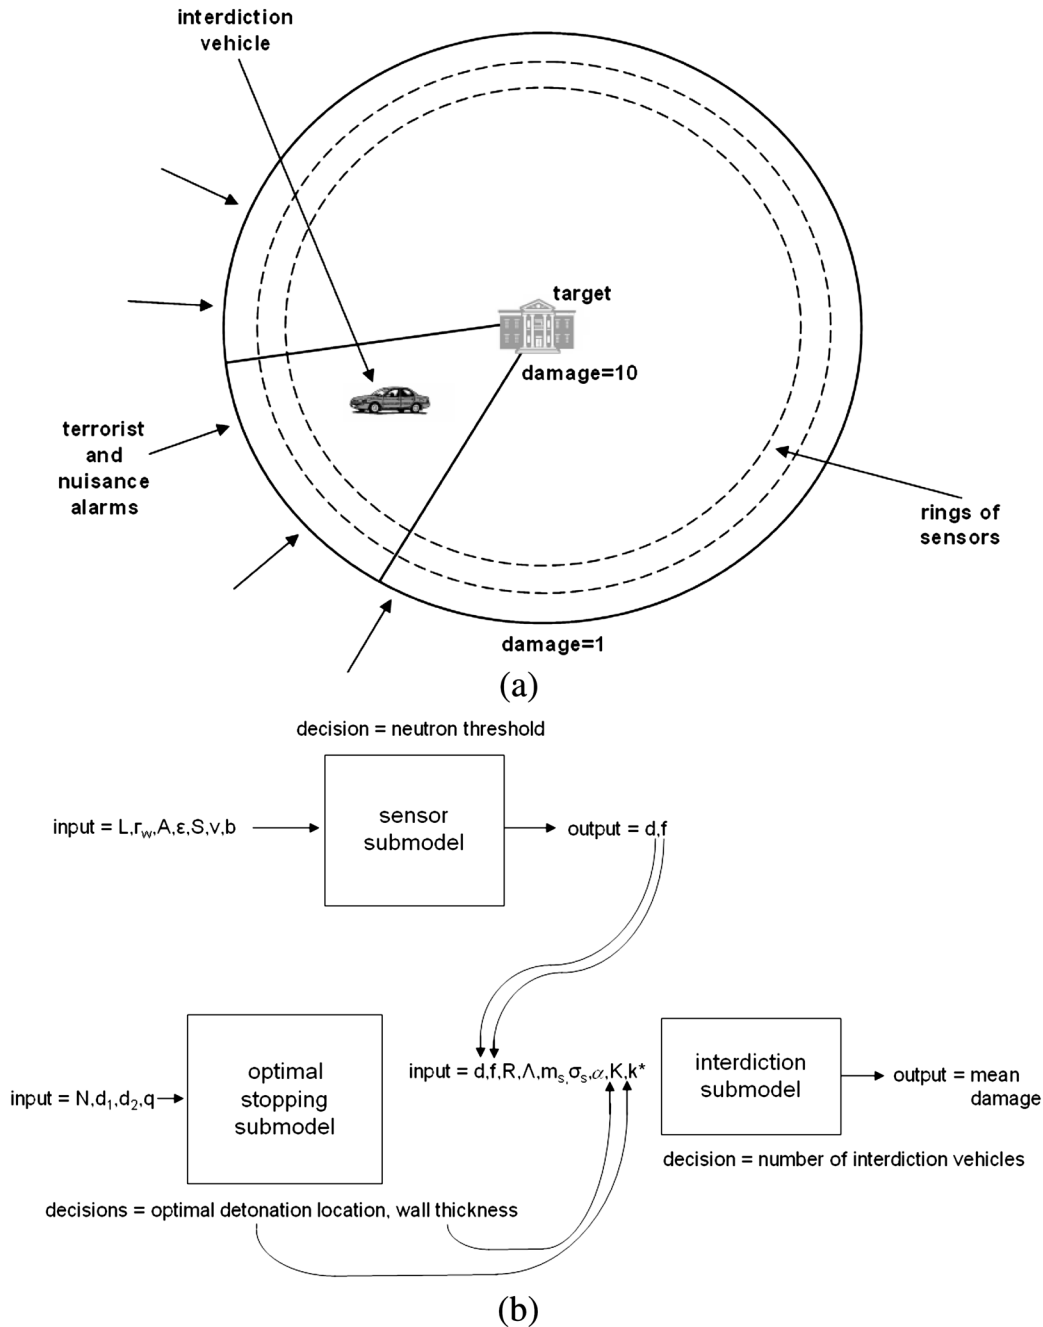


Fig. 1. (a) The model geometry. The government’s optimization problem is to choose the sensor threshold (\bar{n}), the number of rings of sensors (K) and the number of interdictor vehicles (M), one per wedge, to minimize the mean damage inflicted by a terrorist subject to the terrorist’s optimal stopping solution and a budget constraint on rings and wedges. (b) A graphical depiction of the three submodels and their interrelationships. The variables are defined in Table I.

given in Table I. The first submodel considers a vehicle (possibly containing a nuclear weapon) passing through a radiation sensor and determines (in terms of the neutron threshold level) the detection probability and false positive probability, which play a key role in the third submodel.

Stochastic dynamic programming is a field of optimization in which a decision maker makes a set of decisions over time in the face of statistical uncertainty [14]. Optimal stopping problems (Chapter 3.4 in [14]) are an important subset of stochastic dynamic programs in which the decision maker at each stage has the binary choice of whether to continue or to stop. Our second submodel is an optimal stopping problem for a terrorist who is

traveling toward a target through a sequence of radiation sensors. The terrorist does not know the exact detection probability at each sensor and updates this probability in a Bayesian manner (described in more detail below) as he travels toward the target. At each sensor, the terrorist decides whether to detonate immediately or continue traveling toward the target. The output of this submodel is a probability distribution over where, if at all, the terrorist is detected.

Queueing theory is the mathematical study of waiting lines [15]: It considers a stochastic system in which “customers” arrive randomly over time and set of “servers” provide service to these customers, typically one at a time (the service time is a

TABLE I
VALUES FOR THE MODEL PARAMETERS

Parameter	Description	Value	Reference
L	Vehicle length	5 m	
r_w	Distance from vehicle center to sensor	4 m	Text
A	Detector area	0.3 m ²	[9]
ϵ	Efficiency of neutron detector	0.14	[9]
S	Neutron source strength (plutonium)	400k neutrons/sec	[9]
S	Neutron source strength (uranium)	30 neutrons/sec	[9]
v	Vehicle speed past sensor	10 m/sec	Text
b	Mean neutron background rate	50 neutrons/m ² ·sec	[9]
e^μ	Median of background emissions detected	43.08 neutrons	[10], Appendix I
e^σ	Dispersal factor of background emissions detected	1.73	[10], Appendix I
\bar{n}	Neutron threshold level	Decision variable	
N	Network size	5 or 50	Text
K	Wall thickness	Decision variable	
d_1	Maximum damage	10	Text
d_2	Slope of damage	$\frac{9}{N}$	Text
q	Probability bomb is detonated during interdiction	0.5 or 0.9	
k^*	Terrorist's optimal stopping point	Decision variable	
R	Wall radius	50 miles	
Λ	Arrival rate of cars to circle	10 ⁵ /hr	[11]
M	Number of interdiction vehicles	Decision variable	
m_s	Mean on-site service time	30 min	
σ_s	Standard deviation of on-site service time	3 min	
α	Relative speed of interdiction vehicle	1.5	
Y	Total number of sensors	$\pi K(2N - K + 1)$	Text
c_Y	Annual cost per sensor	\$50k	[12], text
c_M	Annual cost per interdiction vehicle	\$850k	[13], text
B	Annual budget	Varies	

random variable) and in a first-come first-served fashion. The goal of queueing theory is to predict the amount of congestion in the system (e.g., the number of customers waiting in line or the amount of time a customer waits in line). Our third submodel is a spatial version of a queueing model, in which customers arrive randomly over time and space. The model consists of a circle, the perimeter of which represents the outer wall of sensors; the sequence of sensors in the optimal stopping submodel represents a terrorist's straight path from the perimeter of the circle toward the target in the circle center. The customers and servers in this queueing model are mobile: the customers are vehicles that trigger a sensor alarm and are moving toward the target, and the servers are interdiction vehicles that chase the customers and upon capturing them, successfully prevent a bomb detonation with a specified probability. The output of this submodel is the expected location of where (if at all) the terrorist would be interdicted. Because the interdiction vehicles also chase non-terrorist vehicles that trigger an alarm, the output of the third submodel depends on both the detection probability and false positive probability from the sensor submodel.

We now describe the three submodels in more detail.

B. The Sensor Submodel

Radiation sensors detect both neutrons and gamma-rays. Gamma-ray technology is in flux: the high gamma-ray alarm

rates of legal goods generated by the non-spectroscopic technology that has been deployed at ports over the last few years [16] precludes its near-term use in this application, where each nuisance alarm requires interdiction by a police or government vehicle. Spectroscopic gamma-ray detectors, which may be capable of distinguishing fissile material from other legal shipments, are coming on line now and over the next several years [17], but for lack of performance data (data on the detection probability, nuisance alarms, and false alarms, i.e., alarms in the absence of a source such as a legal good or a nuclear medical patient), we restrict our attention to neutron detection. Nonetheless, an important by-product of our analysis is the total false positive probability (due to nuisance alarms and false alarms) required for a feasible deployment of either neutron or gamma-ray sensors.

Our sensor model is rather crude by the standards of this journal: we do not delve into the physics, engineering and signal processing details (e.g., [18]) of the sensors (e.g., we do not consider the details of the directionality and the spectrum of background neutrons [19]), and instead we generalize an existing model [9] developed for detecting nuclear warheads, and in Appendix I we estimate several model parameters so that the sensor performance agrees with the published performance [10] of currently-deployed systems [20]. That is, the goal of our sensor submodel is to mimic the Receiver Operating Characteristic (ROC) curve of currently-deployed sensors, not to improve

the operation of sensors, and our model appears to be at the appropriate level of detail given the paucity of operating data of currently-deployed systems.

According to the only published data that we are aware of, the IAEA conducted roadside testing in Europe and found that 163 000 trucks generated no neutron alarms [10] (although it is conceivable that some legal items did not emit enough neutrons to set off the alarm). However, the U.S. has different shipping regulations than Europe and there are several legal items (e.g., AmBe soil density gauge) that emit neutrons. Initial testing suggests that the neutron nuisance alarm rate $\approx 10^{-4}$ at U.S. border crossings [21]. Because of the lack of detailed data on neutron nuisance alarms (e.g., the emissions from the vehicles generating alarms) and because the neutron nuisance alarm rate is very small, we assume there are no emissions from vehicles that do not contain a nuclear weapon; however, we revisit this issue in Section IV. We explicitly model background neutron emissions, much of it from spallation caused by cosmic rays, and emissions from a weapon. We let $L = 5$ m be the length of a vehicle and $v = 10$ m/sec (i.e., 22.4 mph) be the vehicle speed, which is representative of the minimum speed on a typical exit ramp. Hence, the testing time for a vehicle is L/v . Background emissions are modeled in (1) and (3) of [9] as a random variable with mean and variance $A\epsilon bL/v$, where $b = 50$ neutrons/m²·sec is the mean neutron background rate, $A = 0.3$ m² is the detector area (which is comparable to the 0.228 m² area per pillar for currently-deployed equipment [20]), and $\epsilon = 0.14$ is the efficiency of the neutron detector. On page 235 of [9], these authors assume that the shielding around the detector reduces the background in directions other than toward the source by a factor of 10, so that the parameters A and ϵ are the same for the background and for the source, and we make the same assumption about A and ϵ here. Our model differs from the one in [9] in that we attempt to account for the fact that the intensity of cosmic-ray induced neutrons varies with solar magnetic activity [22], [23] (it also varies with altitude and the location within the geomagnetic field, but these remain constant over time for any given location). We do this by assuming that the background emissions detected by the detector, which is denoted by X_2 , is a Poisson random variable (a Poisson random variable has a mean equal to its variance by definition, and takes on discrete, non-negative values) with mean $A\epsilon bL/v$, but where the background rate b is itself a log-normal random variable with median e^μ and dispersal factor e^σ (i.e., $\ln b$ is a normal random variable with mean μ and standard deviation σ). Hence, the log-normal random variable captures the fluctuations of solar magnetic activity over time (typically on longer time scales than the testing time), while for a given level of activity, the Poisson random variable captures the natural stochasticity of emissions over the short time-scale (i.e., during testing). In Appendix I, we state the probability mass function of X_2 (and of X_1 , which is described shortly) and estimate values for the parameters μ and σ from controlled experiments [10].

Turning to the weapon emissions, we let $r_w = 4$ m be the shortest distance from the weapon to the sensor as the vehicle drives past the sensor, where we have in mind a sensor on the side of a single-lane exit ramp that is parallel with the flow of traffic. For simplicity, we assume that there is no negative in-

terference (i.e., blocking of emissions) from other vehicles on either the weapon or background emissions. We also assume that a vehicle with a strong neutron source causes a single alarm as it passes through a sensor, effectively assuming that sensors are sensitive only to vehicles within their designed field of view and insensitive to emissions from vehicles that are several lanes away. We denote the source strength of the weapon by S , which we take to be 400 k neutrons/sec for a plutonium weapon (4 kg of weapons-grade plutonium) and 30 neutrons/sec for a uranium weapon containing 12 kg of weapons-grade uranium; both of these numbers were computed from weapon models of Soviet nuclear warheads in Table III of [9]. The cumulative emissions at the detector due to a stationary source with testing time L/v at a distance d_w is modeled as $(A\epsilon SL)/(4\pi d_w^2 v)$ in (3) of [9]. Because our source is driving at velocity v and the shortest distance between the weapon and the sensor as the vehicles drives past the sensor is r_w , we calculate the cumulative emissions at the detector by integrating over time, where time $t = 0$ corresponds to when the front of the vehicle passes by the sensor, and $t = L/v$ corresponds to when the back of the vehicle passes by the sensor. The weapon is in the middle of the vehicle (both lengthwise and widthwise), so that the weapon achieves the minimum distance r_w from the sensor at time $t = L/2v$ and the distance between the weapon and the detector at time t is $\sqrt{r_w^2 + (vt - (L/2))^2}$. Performing this integration from 0 to L/v , we find that the cumulative emissions at the detector due to the weapon is a Poisson random variable with mean

$$\frac{A\epsilon S}{4\pi} \int_0^{L/v} \frac{dt}{r_w^2 + (vt - \frac{L}{2})^2} = \frac{\tan^{-1}\left(\frac{L}{2r_w}\right) A\epsilon S}{2\pi v r_w}. \quad (1)$$

Because the sum of independent Poisson random variables is also a Poisson random variable, we assume that the emissions from a weaponized container detected by the detector, which we denote by X_1 , is a Poisson random variable with a mean that is the sum of the constant in (1) and $A\epsilon bL/v$, where b is a log-normal random variable with median e^μ and dispersal factor e^σ .

If we let \bar{n} be the neutron threshold above which an alarm is generated, then the detection probability is $d = P(X_1 > \bar{n})$ and the false positive probability is $f = P(X_2 > \bar{n})$. By varying the threshold \bar{n} , we sweep out ROC curves for plutonium and uranium weapons (Fig. 2). We assume that detection at different nodes are statistically independent events, which is partially justified by the fact that neutron emissions are bursty [24], background noise can vary across time and space, and different nodes have different sensors.

C. The Optimal Stopping Submodel

To maintain consistency between the optimal stopping submodel and the interdiction submodel, we adapt the optimal stopping results in [25], which considers a two-dimensional square lattice, to the one-dimensional linear network consisting of a set of $N + 1$ nodes indexed by $k = 0, \dots, N$, where the nodes are street intersections or highway entrance/exit ramps, and the edges are road segments. The target is at node 0 and if the radiation wall is of thickness K , then radiation sensors are deployed




Fig. 2. For a sensor in isolation, the detection probability vs. false positive probability for a plutonium weapon (—), a uranium weapon (---), and a shielded plutonium weapon with a shielding factor of four (| | |).

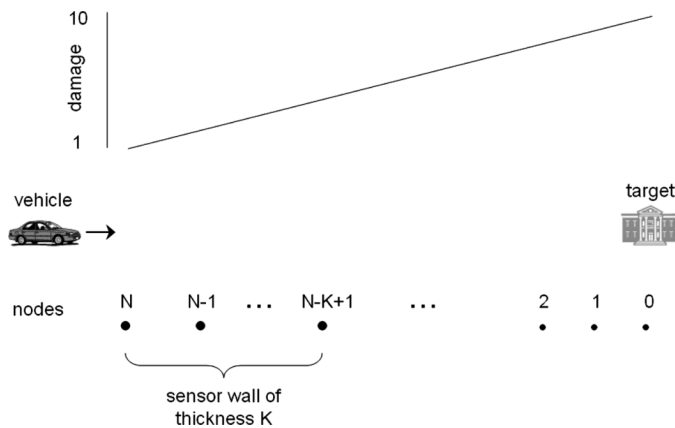


Fig. 3. The optimal stopping problem with the linear damage function.

at nodes $N-K+1, \dots, N$ (Fig. 3). In a two-dimensional lattice, this outer wall topology has been shown to be more effective than a random topology, in which each node contains a sensor with a random probability and fake sensors are placed at all other nodes so that the terrorist cannot observe the exact location of sensors [25].

The terrorist, upon observing the location of the sensors, starts at node N and travels toward the target at node 0. The damage caused by a bomb detonated at node k is assumed to be $d_1 - d_2k$, and we set $d_1 = 10$ and $d_2 = (9/N)$, so that the damage is normalized to be in the 1-to-10 range, which we believe maximizes our results' transparency for policy-makers. The damage can be loosely interpreted as representing a population gradient or the proximity to an important target (e.g., the White House, Times Square) at node 0; because of the uncertainty in the exact nature of the damage function, we carried out an analogous study with a damage function that is exponential, which is perhaps more in line with physical laws [26], and found that the results were qualitatively similar under the linear and exponential damage functions [25].

At each node, the terrorist makes the decision of either detonating the bomb at that node or moving forward based on his perception about the detection probability of the sensor. Because very few legal goods emit neutrons, it would be extremely difficult for a terrorist to probe this network (i.e., to learn the detection probability of a sensor and the timeliness of interdiction). Consequently, the terrorist in our model has no prior information on the detection probability and assumes that interdiction is instantaneous (i.e., there is no time lag between detection and attempted interdiction). As the terrorist travels past a sensor at a node, he updates his perception of the detection probability of the sensors in a Bayesian manner. That is, he has a prior distribution on the detection probability (in other words, detection is a Bernoulli, or 0–1, random variable, but the detection probability is itself a random variable. After passing through a sensor, he uses the information gained on whether or not he was detected to obtain a posterior distribution of the detection probability. In Bayesian problems such as this, it is highly desirable to use a “conjugate” prior distribution (Section 4.2 in [27]), in which the prior and posterior distributions have the same form (but with different parameters). For our problem, in which detection is a 0–1 random variable, the conjugate distribution is the beta distribution (pg. 287 of [27]), and in particular, we assume an uninformative prior, which implies that the detection probability is 0.5 before the terrorist passes through any sensors. It follows that just before passing through node k , the terrorist has successfully passed through nodes $N, \dots, k+1$, and believes that the detection probability is $(1)/(N-k+2)$ (pg. 287 of [27]). In [25], we show that the terrorist can cause slightly more expected damage if he knows the detection probability prior to entering the network. We assume that if the terrorist is detected by the sensor at a certain node, he would try to detonate the bomb at that node, and he would succeed in doing so (before being killed or captured) with a specified probability q , which is known by both the terrorist and the government. Once inside the K layers of the wall, he knows there are no more sensors and that he can travel freely to the target.

The optimal stopping problem is formulated in Appendix II. The solution to the optimal stopping problem is k^* , the node at which the terrorist stops and detonates the bomb. The optimal solution is also given in Appendix II, and is either $k^* = N$ (detonate the bomb before passing through any sensors because he believes that he will cause more damage this way due to his perception that the possibility of being detected and successfully interdicted is too high) or $k^* = 0$ (proceed to the target), depending on the values of the model parameters d_1, d_2, q , and N , and the wall thickness K .

D. The Interdiction Submodel

This submodel derives the damage caused if the terrorist proceeds directly to the target. The city is modeled as a circle of radius R (50 miles), with generic state r , where $r = 0$ denotes the center of the circle, which is the terrorist's target, and $r = R$ is the outer perimeter of the radiation wall. To maintain consistency with the optimal stopping subproblem, the distance between sensors is (R/N) . Hence, if the sensor wall consists of K layers (or rings), then sensors are placed at radii $((N-i+1)R)/(N)$ for $i = 1, \dots, K$ and the total

number of sensors is the sum of the perimeters of these K layers, $\sum_{i=1}^K (2\pi(N-i)R)/(N)$, divided by the distance between sensors, (R/N) , which we denote by $Y = \pi K(2N - K + 1)$.

We assume vehicles arrive according to a Poisson process at rate Λ (i.e., the time between consecutive arrivals is an exponential random variable with mean Λ^{-1}) uniformly around the circle perimeter (i.e., vehicles arrive randomly in time and space) and then travel directly toward the circle center at speed R miles per hr, so that it takes 1 hr to travel from the perimeter to the center. Note that while a speed of 50 miles per hr might not be achievable in the most congested U.S. cities, even during off-peak hours, the critical assumption is not the speed of the vehicle, but that it takes one hour to travel from the edge of the city to the center. We do not model outbound traffic (a terrorist possessing a nuclear weapon in a major city would be unlikely to leave the city with it) or track vehicles after they reach the circle center. In 2001, 1.446 M vehicles paid tolls daily in the New York City area, 85% of which arrived during 7 am–7 pm [11], and so we set $\Lambda = 10^5/\text{hr}$. Because the arrival of a vehicle with a nuclear weapon is an extremely rare event, the amount of congestion in our spatial queue is dictated by the false positive alarms. If the sensor wall is K layers thick, then the Poisson arrival rate of customers at radius r in the queueing system, denoted by λ_r , is $\lambda_{((N-i+1)R)/(N)} = \Lambda(1-f)^{i-1}f$ for $i = 1, \dots, K$.

Following [28], we divide the circle into M equal wedges and assign one interdiction vehicle to each wedge, thereby allowing us to analyze a single wedge, which has customer arrival rates $\lambda_{((N-i+1)R)/(N)}/M$ for $i = 1, \dots, K$, where these arrivals are uniformly distributed on each of the K arcs. The interdiction vehicle travels at rate αR (where $\alpha = 1.5$) and its movement is restricted to be along any ray (i.e., line emanating from the circle center) and any ring (i.e., constant radius path). The interdiction vehicle serves customers in a first-come first-served manner. The interdiction vehicle first chases a customer and follows the rays and rings that minimize the distance a customer travels before being caught. If the customer is located at radius r_c and the interdiction vehicle is located at radius r_s before the chase begins, then the server will chase the customer if and only if the customer can be caught before he reaches his target, i.e., if $r_c > (r_s)/(\alpha)$. Otherwise, the server ignores the customer, who succeeds in reaching the target. Upon catching a customer, the interdiction vehicle performs an on-site service (to determine the source of the radiation emissions, do a background check on the driver and vehicle, etc.) that is a normal random variable with mean m_s (in units of time) and standard deviation σ_s . At the completion of service, the interdiction vehicle returns to its optimal resting location if no other customers are in the wedge; if other customers are in the wedge, then the next catchable customer is pursued. We choose the optimal resting location to be the point in the wedge that minimizes the expected chase time for a terrorist who arrives to an empty system (because the server is idle the great majority of the time in the damage-minimizing solution, we found that this resting location performs slightly better than the location that minimizes the expected chase time for a false-positive customer who arrives to an empty system); for a system with static customers arriving on the perimeter, the first-come first-served policy with this resting

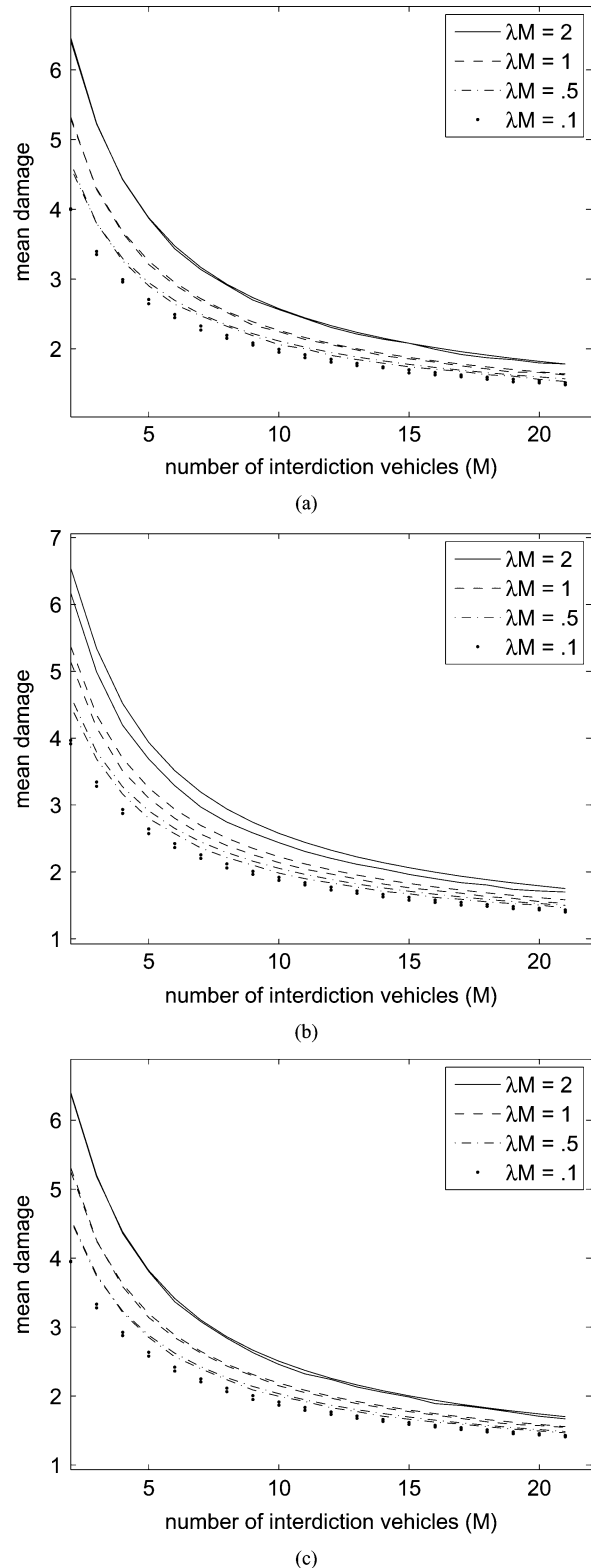


Fig. 4. The accuracy of the mean damage in (A.36). The mean damage (via simulation and (A.36) vs. the number of interdiction vehicles for different values of λ when $q = 0.9$ and (a) $N = 5, K = 1, f = 0.01, d = 0.99$; (b) $N = 50, K = 10, f = 0.01, d = 0.99$; (c) $N = 50, K = 10, f = 0.99, d = 0.99$.

location is shown to minimize the expected time that customers spend in the system as the arrival rate goes to zero [28].

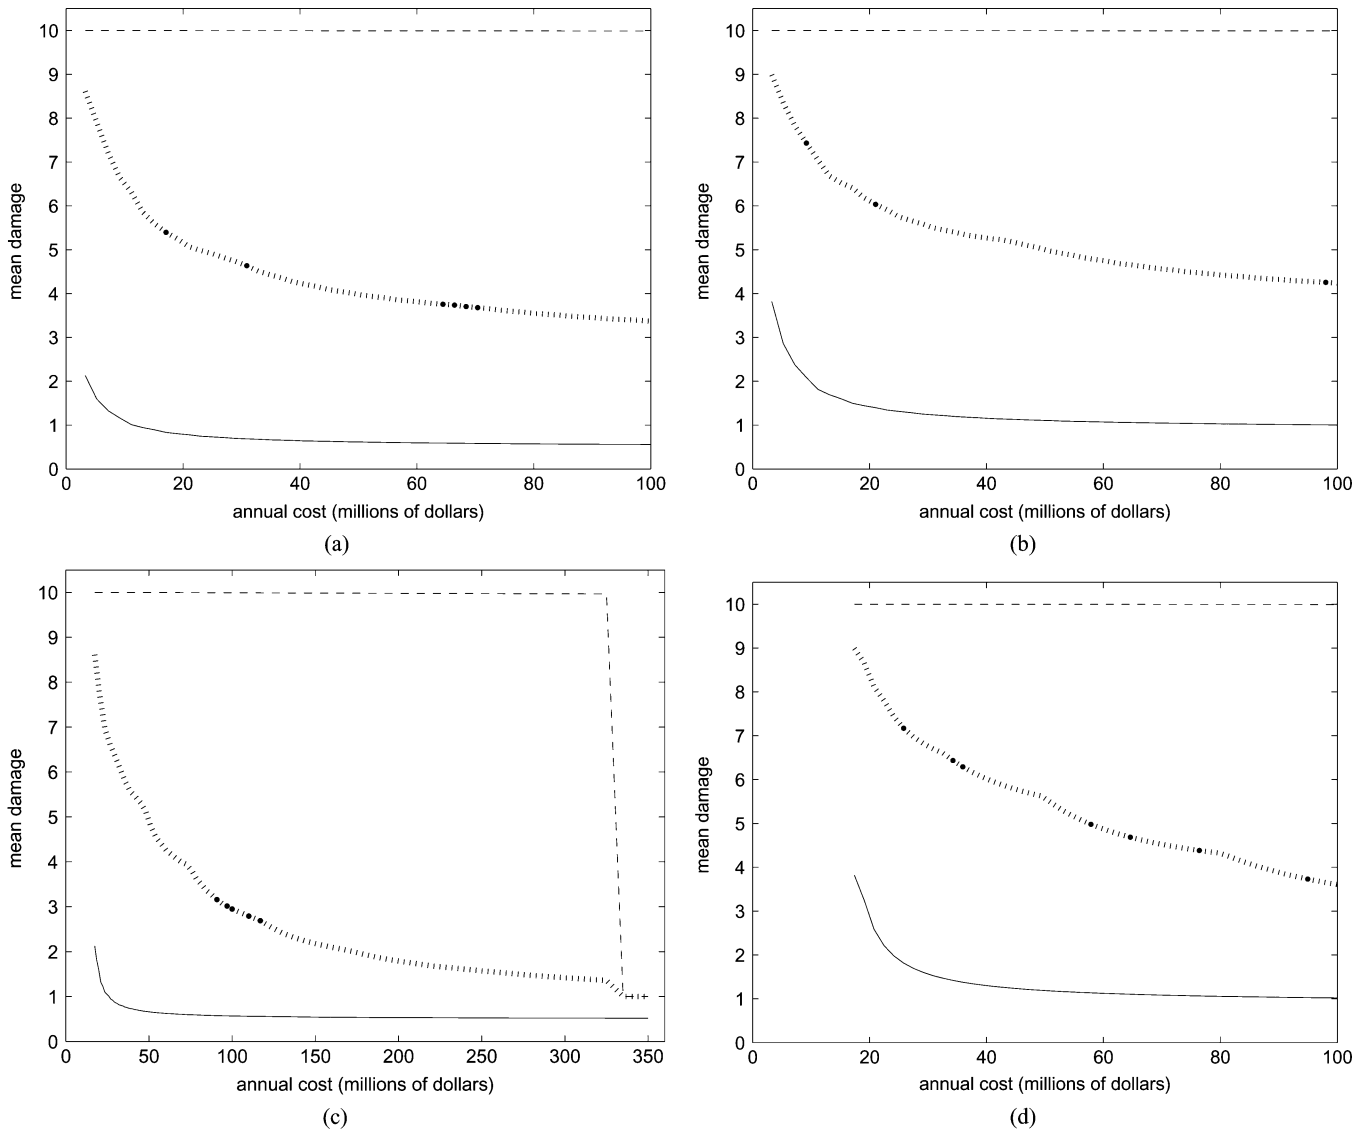


Fig. 5. Mean damage vs. cost curves for plutonium (—), shielded plutonium (·|·|) with a shielding factor of four, and uranium (---) weapons, for (a) small network ($N = 5$) and high interdiction ($q = 0.5$) (b) small network ($N = 5$) and low interdiction ($q = 0.9$) (c) large network ($N = 50$) and high interdiction ($q = 0.5$), and (d) large network ($N = 50$) and low interdiction ($q = 0.9$).

This interdiction model was simulated and compared to the analytical results of an approximating single-server queuing system on a wedge that has a limit of two customers in the system (this model is a generalization of the one in [25]), and the results of the two models were found to be in close agreement (Fig. 4). This allows us to use the approximate analytical results (Appendix III) in place of the simulation model, which enables optimization of the government's three decision variables. The objective function is to minimize the mean damage generated by a terrorist vehicle, which can be computed using the optimal stopping and interdiction submodels: if the terrorist chooses $k^* = 0$, then he is detected at node i with probability $(1 - d)^{N-i}d$ for $i = N - K + 1, \dots, N$, and makes it to the target with probability $1 - \sum_{i=N-K+1}^N (1 - d)^{N-i}d$, and the interdiction model computes the ultimate fate (where and if they are caught) of a terrorist detected at node $i = N - K + 1, \dots, N$ (Appendix III). If the terrorist chooses $k^* = N$, then the damage equals $d_1 - d_2N = 1$, regardless of interdiction resources.

III. RESULTS

We choose the wall thickness (K), the neutron threshold level (\bar{n}), and the number of interdiction vehicles (M) to minimize the expected damage subject to the terrorist choosing the optimal stopping solution k^* and subject to the budget constraint $c_Y Y + c_M M \leq B$, where c_Y is the annual cost per sensor, c_M is the annual cost per interdiction vehicle, and B is the annual budget. The cost of a radiation portal monitor is approximately \$100 k [12]. Adding installation, which is nearly as expensive as the equipment, and maintenance and considering a five-year lifetime, we set $c_Y = \$50$ k/yr. These sensors are attended remotely by the drivers of the interdiction vehicles. An interdiction vehicle requires eight workers (two people for three shifts plus a backup shift) at \$100 k/yr plus two \$100 k vehicles (one backup) [13] prorated over a five-year lifetime, for a total of $c_M = \$850$ k/yr. By solving this problem for various budget levels, we generate expected damage vs. cost curves

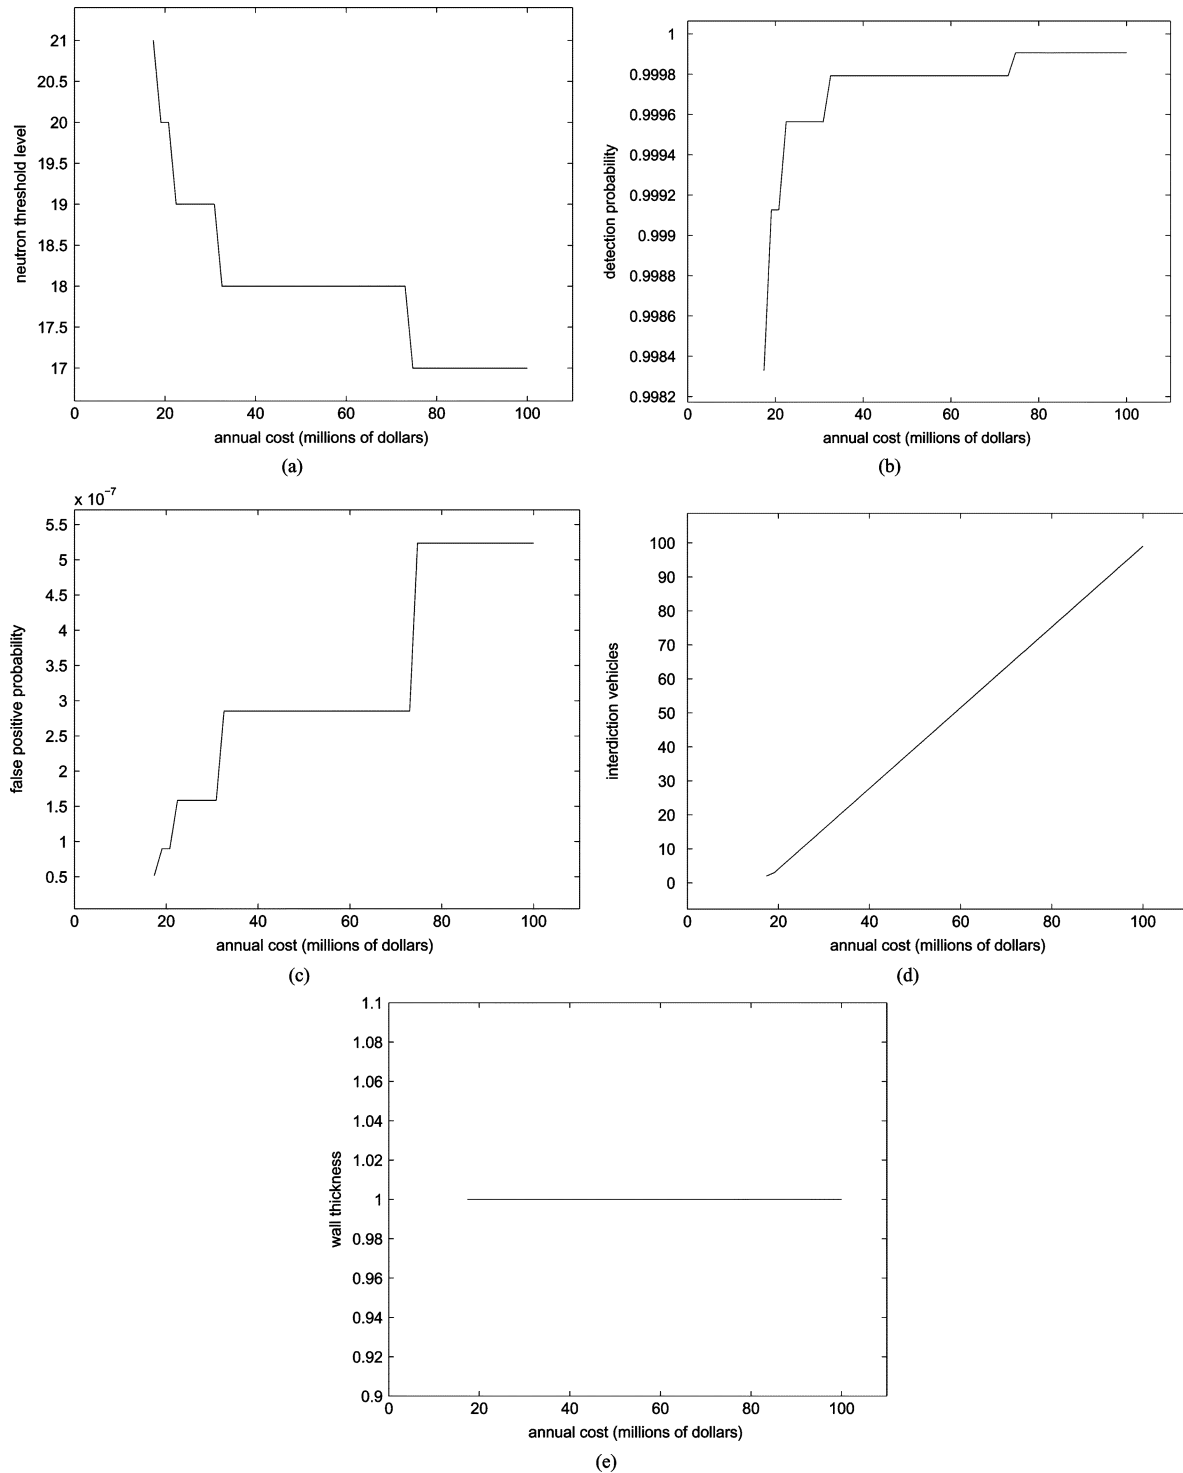


Fig. 6. The optimal solution vs. the annual cost (i.e., budget) for the ($N = 50, q = 0.9$) plutonium case. (a) The optimal neutron threshold level (\bar{n}^*); (b) the detection probability in (A.38); (c) the false positive probability in (A.39); (d) the optimal number of interdiction vehicles (M^*); (e) the optimal wall thickness (K^*).

for both plutonium and uranium weapons in four scenarios: combinations of small ($N = 5$) and large ($N = 50$) networks and high ($q = 0.5$) and low ($q = 0.9$) interdiction (Fig. 5). The small network represents a city with several chokepoints (e.g., tunnels, bridges) such as New York City and the large network represents a city with many highways such as Los Angeles.

The terrorist cannot be deterred from proceeding directly to the target in three of the four scenarios. In the ($N = 50, q = 0.5$) scenario, the terrorist detonates the bomb before passing through any sensors when the wall thickness $K \geq 29$ [Fig. 5(c)], which requires an annual investment of \$328 M.

Because a sensor can achieve a high detection probability and a low false positive probability for a plutonium weapon (Fig. 2),

it is optimal in the plutonium case to have a thin wall (i.e., $K^* = 1$), even for large networks and reasonably large budgets [Fig. 6(e)]. As the budget increases, the optimal neutron threshold level (\bar{n}^*) is nonincreasing, the false positive probability increases with the budget, the optimal number of interdiction vehicles (M^*) increases linearly with the budget so as to maintain a low level of congestion where vehicles are idle approximately 95% of the time (Fig. 6), and the expected damage decreases but with decreasing returns (Fig. 5). Assuming low interdiction ($q = 0.9$), a single-layer wall and approximately 10 interdiction vehicles can reduce the mean damage to approximately 2 (on the 1-to-10 scale) for an annual cost of approximately \$10 M for a small network and \$25 M for a large network.

In contrast, a sensor's detection probability and false positive probability are approximately equal for a uranium weapon (Fig. 2), and the detection-interdiction system has virtually no impact: e.g., in the ($N = 50, q = 0.9$) case, the mean damage is 9.997 when the annual budget is \$55 M ($K^* = 1, M^* = 45$). Because the plutonium and uranium weapons generate extreme cases that are very easy and difficult to detect, respectively, and because terrorists can shield a weapon's emissions by surrounding it with a heavy metal such as lead or tungsten, we also consider an intermediate case in which the neutron source strength $S = 100000$ neutrons/sec, which represents a plutonium weapon with an additional shielding factor of four. In this case, it is optimal to add additional rings of sensors, and $K^* = 2$ in the small network and $K^* = 3$ in the large network when the annual budget is \$100 M; in the large network, three rings achieve a detection probability of 0.84 although each sensor's detection probability is approximately 0.46 (Fig. 7). The additional rings are slightly more effective in the large networks because they are farther from the target than in the small network (the curves for the smaller networks flatten out faster than the larger networks in Fig. 5). There are jumps in Fig. 7 because all of the decision variables are discrete, however the nonmonotonic nature of the figure is a result of the tradeoffs between having more vehicles versus having more rings and between congestion and detection. Once it is optimal to add another ring, the number of interdiction vehicles decreases in order to satisfy the budget constraint. With fewer interdiction vehicles and more rings, congestion in the system increases. However, with an additional ring, more terrorists can be detected with a lower value of d and thus the government can afford to decrease \bar{n} , which reduces congestion. It appears that a moderate amount of shielding makes it difficult to mitigate the damage generated by an attack. More generally, in a small network, there is a phase transition in which a drop in neutron emissions by 20 k neutrons/sec (from 110 k to 90 k when $q = 0.5$ and from 130 k to 110 k when $q = 0.9$) causes an easy-to-detect situation to change into a difficult-to-detect situation [Fig. 8(a)]. For large networks, the transition is more gradual, occurring in the 70 k–130 k neutrons/sec range, because additional rings of sensors are more helpful [Fig. 8(b)].

IV. DISCUSSION

Although the optimal stopping solution suggests that a wall more than 29 sensors thick in the large network would deter a terrorist from passing through a sensor if the probability that the bomb can be detonated during interdiction is 0.5, this result

should be viewed with some skepticism, even aside from the exorbitant cost of such a dense sensor network. It is very difficult to predict terrorist behavior, and a terrorist that can detonate remotely from the driver's seat is likely to view the likelihood of detonation during interdiction as being near 1. Hence, combined with the possibility that terrorists are risk-seeking rather than risk-neutral, we believe it is prudent to assume that a terrorist will proceed directly to the target, although the possibility of deploying a dense set of fake (i.e., inoperable) sensors in addition to the real sensors should be investigated.

The spatial interdiction model is merely a caricature of an actual highway system, and our main contribution may be in the framing of the problem rather than in the numerical results. Some cities (e.g., New York City) are not laid out with concentric highway grids where arrivals to the city perimeter are uniformly distributed. More refined insights would require the modeling of a specific city's highway structure or possibly the modeling of highways by a percolation process [29], which would be much more difficult to analyze. They would also require incorporating other operational issues, such as non-homogeneous traffic rates to capture rush hour, and the cooperation of interdiction vehicles in adjacent wedges. Moreover, the sensitivity and specificity data we use is from the late 1990s and may not be representative of the equipment that is currently in use, whose performance metrics are proprietary; this precludes the direct use of Figs. 5 and 8 by the U.S. Government or a nuclear terrorist.

Nonetheless, our qualitative results are likely to be robust: A detection-interdiction system with a single layer of sensors and 10 dedicated, but mostly idle, interdiction vehicles can mitigate the damage caused by an unshielded or lightly-shielded weapon made of plutonium, but not uranium. Although our results echo the performance of existing detection systems at ports [8], [30], our setting is more difficult because relative to ports, on a highway vehicles will typically move more quickly past the detectors and at a greater distance from the detectors. A system with several layers of sensors can help offset moderate shielding of a plutonium weapon as well as make it more difficult for a terrorist to somehow bypass (e.g., using off-road transport) the outermost layer of sensors.

Our results only pertain to the detection of neutron emissions; spectroscopic gamma-ray detectors, which may be capable of distinguishing fissile material from other legal shipments, are in the process of being deployed [17]. This new technology could be particularly helpful in detecting 2.614 MeV emissions, which are emitted from former Soviet nuclear warheads containing ^{232}U from reprocessed reactor fuel (page 238 in [9]). The false positive probability for a gamma-ray detector would need to be no more than 10^{-4} to mitigate the damage (e.g., using $\Lambda = 10^5/\text{hr}$ in the solid curve in Fig. 4(a), 10 interdiction vehicles achieves a mean damage of 2.7 with a false positive probability of 2×10^{-5}) with a budget of tens of millions of dollars; interestingly, 10^{-4} is also the estimated order of magnitude of the neutron nuisance alarm rate in the U.S. [21], suggesting that the optimal neutron threshold level would be set so that the nuisance alarm rate and the false alarm rate (i.e., an alarm in the absence of a source) are both $\approx 10^{-4}$. Non-spectroscopic gamma-ray detectors recently generated false alarms at the rate

of 0.025 at the Port of New York and New Jersey, although Pacific Northwest Laboratories advised the port managers that the alarm rate would be 10-fold less [16]. The Department of Homeland Security predicts that the nuisance alarm rate should decrease by a factor of ≈ 50 with the implementation of spectroscopic gamma-ray detectors [17], suggesting a false alarm rate of 5×10^{-4} , which is approaching the required order-of-magnitude. This level of specificity appears unlikely for detecting radiological dispersal devices, which could generate similar emission profiles as legal items, unless the great majority of vehicles that ship legal radiological items are pre-registered and are not pursued if they trigger an alarm (which itself opens up a vulnerability for terrorists to exploit by stealing one of these vehicles). To adapt our results for gamma detection of radiological dispersal devices, it might be more appropriate to assume that the terrorist has probed the network with legal items that emit gamma rays, and consequently knows the detection probability and the time from detection until interdiction; this last quantity could be incorporated into the optimal stopping problem through a space (i.e., number of nodes) lag between detection and interdiction. If the terrorist can probe the network, it is more difficult to deter him from pursuing his target because he sometimes erroneously plays too boldly in the Bayesian setting [25].

Detection systems are a waste of money if they are not reinforced by a strong interdiction system. A recent detection-interdiction study of pedestrian suicide-bombers found that the damage cannot be significantly reduced because there is not enough time to effectively interdict [31]. In our model, effective interdiction can be achieved because the terrorist is detected tens of miles from his target, and interdiction resources are not highly dependent on the size of the network: it takes approximately \$10 M per year to effectively interdict a plutonium weapon in a small or large network. In contrast, detection resources can vary greatly: the annual cost of a single layer of sensors is \$1.6 M for a small network and \$15.7 M for a large network.

V. CONCLUSION

Taken together, our results suggest that a detection-interdiction system with a single layer of sensors (using currently-deployed technology) and 10–20 dedicated interdiction vehicles could mitigate the damage from an unshielded or lightly-shielded plutonium weapon, but not a uranium weapon or a radiological dispersal device. The annual cost per city, which ranges from several millions of dollars to several tens of millions of dollars, is not incommensurate with the budget of the Domestic Nuclear Detection Office, which has \$193 M in research and development funding for fiscal year 2006 [32]. Additional layers of sensors could help offset some shielding, but at considerable expense. To the extent that a terrorist who can get an assembled nuclear weapon to the outskirts of a U.S. city is likely to shield the weapon if he suspects that a sensor network is in place, the efficacy of an overt system using current technology is questionable. Moreover, although a single-layer system appears to work for a lightly-shielded plutonium weapon in theory, several daunting challenges (aside from developing highly-specific spectroscopic gamma-ray detectors) remain: to design a single-layer system tens of miles

from a city center that cannot be easily bypassed by a vehicle, to design a sensor that can operate in a harsh and insecure environment, and to develop the supporting communications infrastructure so that, e.g., vehicles can be tracked after they set off an alarm.

APPENDIX I THE SENSOR SUBMODEL

The sensor model is described in detail in the main text. In this section, we state the probability mass functions (pmf) for X_1 and X_2 , and then estimate the values for the parameters μ and σ . The probability mass function of a Poisson random variable with mean λ is

$$\frac{\lambda^x e^{-\lambda}}{x!} \quad \text{for } x = 0, 1, \dots, \quad (\text{A.1})$$

and the probability density function of a lognormal random variable is

$$\frac{\exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma x} \quad \text{for } x \geq 0. \quad (\text{A.2})$$

Hence, the pmf for X_2 is

$$p_2(x) = \int_0^\infty \frac{(A\epsilon y L/v)^x e^{-A\epsilon y L/v}}{x!} \times \frac{\exp\left(-\frac{(\ln y - \mu)^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma y} dy \quad \text{for } x = 0, 1, \dots \quad (\text{A.3})$$

By similar reasoning, the pmf of X_1 is

$$p_1(x) = \int_0^\infty \frac{\left(\frac{\tan^{-1}\left(\frac{L}{2r_w}\right)A\epsilon S}{2\pi v r_w} + \frac{A\epsilon y L}{v}\right)^x}{x!} \times \exp\left(-\left[\frac{\tan^{-1}\left(\frac{L}{2r_w}\right)A\epsilon S}{2\pi v r_w} + \frac{A\epsilon y L}{v}\right]\right) \times \frac{\exp\left(-\frac{(\ln y - \mu)^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma y} dy \quad \text{for } x = 0, 1, \dots \quad (\text{A.4})$$

Now we estimate the values of the parameters μ and σ that appear in (A.3)–(A.4). The mean of a lognormal random variable (and hence of the background rate b) is $e^{\mu + (\sigma^2)/2}$. After we determine σ , which is done below, we find μ so that the mean of b equals 50 neutrons/m²·sec:

$$e^{\mu + \frac{\sigma^2}{2}} = 50. \quad (\text{A.5})$$

To determine σ , we turn to extensive controlled experiments [10]. Parameter values for A , ϵ and b are as in Table I of the main text. When the detection time was $\tau = 10$ sec, the distance between the stationary source and the detector was $d_w = 2$ meters, and the source strength was $S = 20$ k

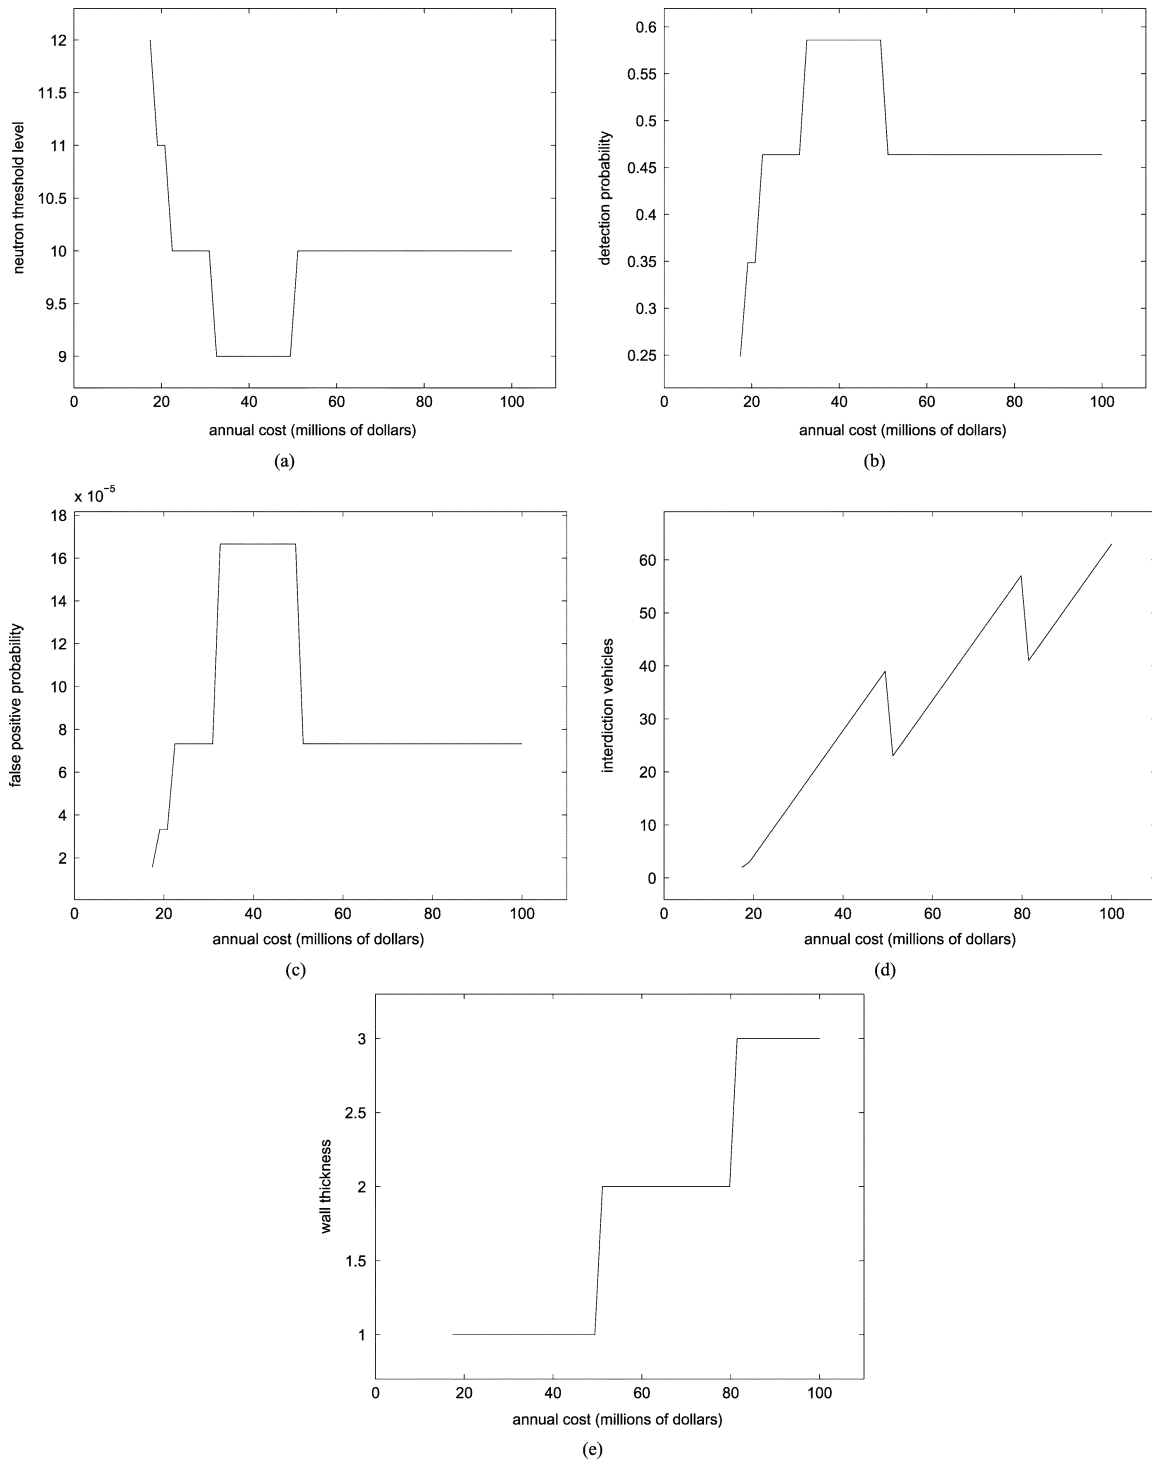


Fig. 7. The optimal solution vs. the annual cost for the ($N = 50, q = 0.9$) shielded plutonium case with a shielding factor of four. (a) The optimal neutron threshold level (\bar{n}^*); (b) the detection probability in (A.38); (c) the false positive probability in (A.39); (d) the optimal number of interdiction vehicles (M^*); (e) the optimal wall thickness (K^*).

neutrons/sec, the false negative probability and false positive probability in these experiments were 10^{-3} and 10^{-4} , respectively [10]. Replacing L/v by τ in (A.3)–(A.4) and replacing $(\tan^{-1}((L)/(2r_w))A\epsilon S)/(2\pi v r_w)$ by $A\epsilon S\tau/4\pi d_w^2$ in (A.4) because these experiments involved stationary sources, and substituting all of these values into

$$P(X_1 \leq \bar{n}) = 10^{-3}, \quad (\text{A.6})$$

$$P(X_2 > \bar{n}) = 10^{-4}, \quad (\text{A.7})$$

yields two equations for two unknowns, σ and the neutron threshold limit \bar{n} in [10]. Solving (A.6)–(A.7) for the two unknowns yields $\bar{n} = 140$ and the dispersal factor $e^\sigma = 1.73$. Substituting σ into (A.5) gives the median $e^\mu = 43.08$ neutrons/m²·sec.

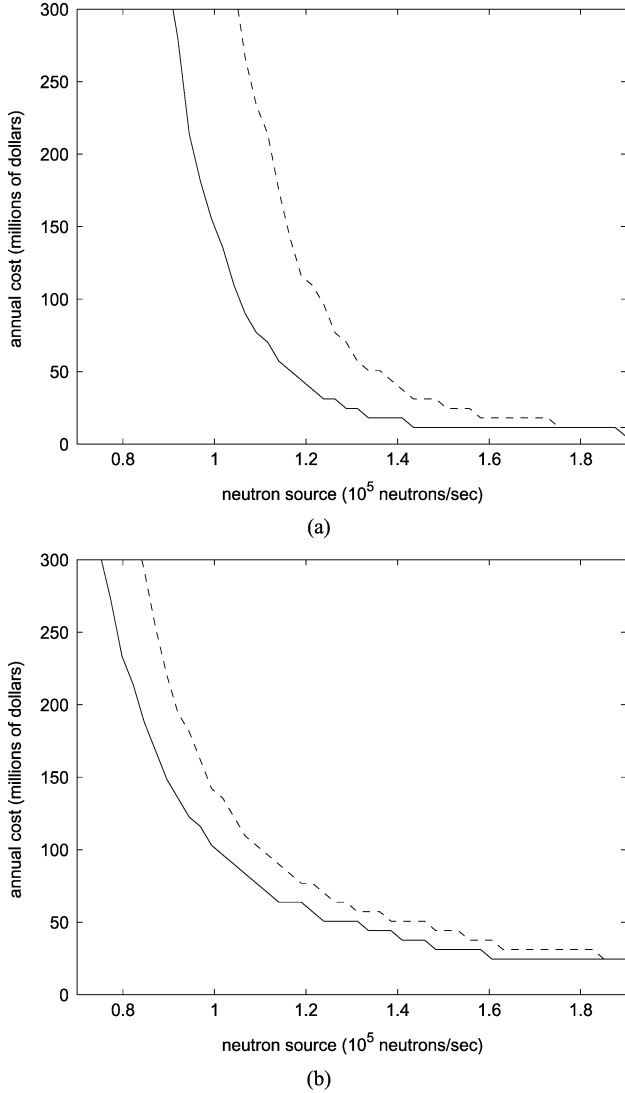


Fig. 8. The minimum cost necessary to maintain the mean damage below 3 (on the 1-to-10 scale) as a function of the neutron source strength, where, for reference, 400,000 neutrons/sec corresponds to an unshielded plutonium weapon, for $q = 0.5$ (—) and $q = 0.9$ (---), for (a) small network ($N = 5$) and (b) large network ($N = 50$).

APPENDIX II

THE OPTIMAL STOPPING SUBMODEL

We formulate the optimal stopping subproblem and provide its solution. In our model, the terrorist makes his detonate vs. proceed decision just before he passes through the sensor at each node. Suppose the terrorist manages to arrive at (but not yet pass through) node k without being caught, where $k \geq N - K + 1$ (so that there is a sensor at this node). Now he has two choices. He can either detonate the bomb at this node or move to node $k - 1$ in an attempt to increase the damage inflicted. Because the event that the terrorist gets detected or not is a Bernoulli random variable at each node, the beta-binomial conjugate pair is a natural choice for modeling the terrorist's updating process for the detection probability at a node. Furthermore, we assume that the terrorist has no prior information about a sensor's detection

probability, and uses the uniform distribution as an uninformative prior to represent his initial perception. Just before passing through node k , the terrorist has successfully passed through nodes $N, \dots, k + 1$. By the above assumptions, the terrorist believes that he will bypass detection at the upcoming node with probability $(N - k + 1)/(N - k + 2)$ and will be detected with probability $(1)/(N - k + 2)$ [27]. If he detonates the bomb before passing through state k , it causes damage $d_1 - d_2k$. If he instead proceeds through state k , he avoids detection with probability $(N - k + 1)/(N - k + 2)$, in which case he travels to node $k - 1$. If the terrorist is detected as he passes through node k (which occurs with probability $(1)/(N - k + 2)$), then with probability q he detonates the bomb before being killed or captured and causes $d_1 - d_2k$ in damage, and with probability $1 - q$ he causes no damage. Because the terrorist can observe the wall of sensors, once he passes through node $N - k + 1$, he proceeds directly to the target at node 0. Taken together, if $V(k)$ is the optimal value function (i.e., maximum expected damage if at node k) then the terrorist's optimal stopping subproblem can be formulated as

$$V(k) = \max \left\{ d_1 - d_2k, \frac{q}{N - k + 2}(d_1 - d_2k) + \frac{N - k + 1}{N - k + 2}V(k - 1) \right\} \quad (\text{A.8})$$

$$\text{for } k = N, \dots, N - K + 1,$$

$$V(N - K) = V(0) = d_1. \quad (\text{A.9})$$

The solution to the optimal stopping problem (A.8)–(A.9), which is derived in [25], is

$$k^* = \begin{cases} 0 & \text{if } d_1 - d_2N \leq \frac{q(d_1 - d_2i)}{\sum_{i=N-K+1}^N \frac{q(d_1 - d_2i)}{(N-i+1)(N-i+2)} + \frac{d_1}{K+1}} \\ N & \text{if } d_1 - d_2N > \frac{q(d_1 - d_2i)}{\sum_{i=N-K+1}^N \frac{q(d_1 - d_2i)}{(N-i+1)(N-i+2)} + \frac{d_1}{K+1}} \end{cases} \quad (\text{A.10})$$

APPENDIX III

THE INTERDICTION SUBMODEL

We compute the approximate expected damage assuming the terrorist chooses $k^* = 0$. The wall is of thickness K and there are M interdiction vehicles, so that the single-server wedge under consideration spans the angles $[0, (2\pi)/(M)]$. For $i = 1, \dots, K$, nuisance arrivals to the wedge occur according to a Poisson process at rate

$$\lambda_i = \frac{\Lambda(1-f)^{i-1}f}{M} \quad (\text{A.11})$$

at radius

$$r_i = \frac{(N - i + 1)R}{N}, \quad (\text{A.12})$$

and the individual arrivals to each constant-radius arc segment are uniformly distributed along $[0, 2\pi/M]$. The analysis has three main steps: calculating the optimal resting location of the server, estimating congestion from nuisance customers via a $M/M/1/2$ queue (i.e., a single-server queue with Poisson arrivals, exponential service times and space for two customers [15]), and computing the damage inflicted by a terrorist who arrives to this queue in steady state. The analysis relies heavily on [25], which performs the same analysis for the case in which all nuisance customers arrive at the same radius.

The optimal resting location of the interdiction vehicle has angle (π/M) by symmetry and uniformity, and we denote its radius by r^* . If all nuisance arrivals occurred at the generic radius r , then we know from (29) in [25] that an estimate for the optimal resting radius is $(r)/(1 + \pi/\alpha M)$. We numerically investigated two possible schemes for the optimal resting location: one location minimizes the mean chase time for a nuisance customer arriving to an empty system and the other location minimizes the mean chase time for a terrorist arriving to an empty system. These are not equivalent because, if detected, a terrorist arrives to radius r_i with probability proportional to $(1-d)^{i-1}d$, whereas a nuisance customer arrives with probability proportional to $(1-f)^{i-1}f$. Intuitively, the terrorist-based location should lead to less damage if congestion in the queue is not too large, and indeed we found that the interdiction vehicles are idle the great majority of time in the optimal solution, causing the terrorist-based location to be superior (although the difference was very small). Hence, because a terrorist is detected at radius r_i with probability $(1-d)^{i-1}d$ for $i = 1, \dots, K$, we approximate the optimal resting radius by the weighted average

$$r^* = \sum_{i=1}^K \frac{(1-d)^{i-1}}{\sum_{j=1}^K (1-d)^{j-1}} \left(\frac{r_i}{1 + \frac{\pi}{\alpha M}} \right). \quad (\text{A.13})$$

As in [25], we estimate the congestion in a wedge due to nuisance arrivals by a $M/M/1/2$ queue with renegeing, in which customers arriving to find two customers already in the system do not enter the queue (because they cannot be caught by the server) and customers arriving to find one customer already in the system renege (i.e., depart from the queue before receiving service) after an exponential amount of time (again, because they cannot be caught by the server). Due to the difficulty of computing the stationary distribution of a multiclass queue of this type (with one class for each of the K arrival radii), we force-fit the problem into a single-class model that has three parameters: the total arrival rate λ , the service rate μ and the renegeing rate γ . Once we determine these three parameters, the steady-state distribution that there are zero, one or two customers in this queue is given by [25], respectively,

$$\begin{aligned} p_0 &= \frac{1}{1 + \frac{\lambda}{\mu} + \frac{\lambda^2}{\mu(\mu+\gamma)}}, \\ p_1 &= \frac{\lambda/\mu}{1 + \frac{\lambda}{\mu} + \frac{\lambda^2}{\mu(\mu+\gamma)}}, \\ p_2 &= \frac{\frac{\lambda^2}{\mu(\mu+\gamma)}}{1 + \frac{\lambda}{\mu} + \frac{\lambda^2}{\mu(\mu+\gamma)}}. \end{aligned} \quad (\text{A.14})$$

The total arrival rate of nuisance customers to the wedge is

$$\lambda = \sum_{i=1}^K \lambda_i. \quad (\text{A.15})$$

The service time includes the chase time and on-site service, and hence

$$\mu^{-1} = m_t + m_s, \quad (\text{A.16})$$

where m_t is the mean chase time. As in [25], we compute m_t using a fixed-point analysis. Let p_n be the fraction of customers arriving to find one customer in the system who are eventually caught. We set $p_n = 1 - p_r$, where the renegeing probability p_r is computed in (A.30). Let t_i^e be the mean time it takes the server to catch a customer arriving to radius r_i if the server is idling at his resting location at the time of arrival, and let t_i^b be the mean time it takes the server to catch a customer arriving to radius r_i (who does not renege) when the server is busy at the time of arrival. By (46) of [25], the expected pre-capture distance traveled by a customer arriving to an empty system at radius r_i is

$$d_i^e = \begin{cases} r_i - \frac{M}{\pi}(\alpha r_i - r^*) \ln \left(1 + \frac{\pi}{(\alpha-1)M} \right) & \text{if } r_i \leq r^*; \\ r_i + \frac{\alpha M(r_i - r^*)^2(\alpha+2)}{2\pi r^*(\alpha+1)} - \frac{\alpha M r_i (r_i - r^*)}{\pi r^*} & \text{if } r^* < r_i < r^* \left(1 + \frac{\pi}{\alpha M} \right); \\ \frac{r_i - r^* (1 - \frac{\pi}{2M})}{\alpha+1} & \text{if } r_i \geq r^* \left(1 + \frac{\pi}{\alpha M} \right). \end{cases} \quad (\text{A.17})$$

Because customers travel at velocity R , the chase time in an empty system is

$$t_i^e = \frac{d_i^e}{R}. \quad (\text{A.18})$$

We assume

$$t_i^b = \frac{d_i^b}{R}, \quad (\text{A.19})$$

where d_i^b , which is the distance traveled during the chase by a non-renegeing customer arriving at radius r_i , is calculated in (A.35) when we estimate the damage inflicted by a terrorist. If we define

$$\bar{t}_e = \sum_{i=1}^K \frac{\lambda_i}{\lambda} t_i^e, \quad \bar{t}_b = \sum_{i=1}^K \frac{\lambda_i}{\lambda} t_i^b, \quad (\text{A.20})$$

then m_t satisfies

$$m_t = \frac{p_0}{p_0 + p_n p_1} \bar{t}_e + \frac{p_n p_1}{p_0 + p_n p_1} \bar{t}_b, \quad (\text{A.21})$$

$$\begin{aligned} &= \frac{1}{1 + p_n \lambda (m_s + m_t)} \bar{t}_e \\ &\quad + \frac{p_n \lambda (m_s + m_t)}{1 + p_n \lambda (m_s + m_t)} \bar{t}_b. \end{aligned} \quad (\text{A.22})$$

Where (A.22) follows from (A.14) and (A.16). Solving (A.22) for m_t yields

$$m_t = \frac{-(1 + \lambda p_n m_s - \lambda p_n \bar{t}_b)}{2p_n \lambda} + \frac{\sqrt{(1 + \lambda p_n m_s - \lambda p_n \bar{t}_b)^2 + 4p_n \lambda (\bar{t}_e + p_n \lambda m_s \bar{t}_b)}}{2p_n \lambda} \quad (\text{A.23})$$

and substitution of this quantity into (A.16) gives the service rate μ for the $M/M/1/2$ system (in terms of p_n and d_i^b , which are calculated later). In [25], we estimate the renegeing rate γ by equating $(\gamma)/(\gamma + \mu)$, which is the renegeing probability of a customer who arrives to find one other customer in the $M/M/1/2$, to $P(T_1 < T_2)$, where the random variables T_1 and T_2 are, respectively, approximate representations of the time until an arriving customer is uncatchable and the residual service time of the customer currently in service. We take the same general approach here, but need to account for the fact that customers arrive at different radial locations. As in [25], we assume for simplicity that T_2 is exponential. To compute its mean, we compute the weighted average of the mean residual service time, where the weights are according to the arrival rates at the various radial locations. We assume the time to chase a customer who arrives at a given radial location is uniformly distributed between the chase time when the customer and server have the same angular location and the chase time when their angular locations differ by (π/M) ; while the chase times are uniformly distributed in [25], this is an approximation here. By (27) in [25], we assume the chase time for a customer arriving to radius r_i is $U[t_i^{\min}, t_i^{\max}]$, where

$$t_i^{\min} = \begin{cases} \frac{r_i - r^*}{(\alpha+1)R} & \text{if } r_i \geq r^*; \\ \frac{r^* - r_i}{(\alpha-1)R} & \text{if } r_i < r^*, \end{cases} \quad (\text{A.24})$$

$$t_i^{\max} = \begin{cases} \frac{r_i - r^*(1 - \frac{\pi}{M})}{(\alpha+1)R} & \text{if } r_i - \frac{r^* \pi}{\alpha M} \geq r^*; \\ \frac{r^* - r_i(1 - \frac{\pi}{M})}{(\alpha-1 + \frac{\pi}{M})R} & \text{if } r_i - \frac{r^* \pi}{\alpha M} < r^*. \end{cases} \quad (\text{A.25})$$

By (39) in [25], the mean residual service time for a customer who arrived at radial location r_i is $m_{r,i} = (\sigma_s^2 + 1/12(t_i^{\max} - t_i^{\min})^2 + (t_i^e + m_s)^2)/(2(t_i^e + m_s))$, and we assume that T_2 is exponentially distributed with mean m_r (and rate $\mu_r = m_r^{-1}$), where

$$m_r = \mu_r^{-1} = \sum_{i=1}^K \frac{\lambda_i}{\lambda} m_{r,i}. \quad (\text{A.26})$$

The time until an arriving customer is uncatchable depends on where the customer arrives, and we let T_{1i} be the time until a customer is uncatchable given that he arrived at radius r_i . For mathematical simplicity, we assume that T_{1i} is uniformly distributed. To determine the parameters of this random variable, we need to estimate the range of locations for the server, who we assume is serving a customer that arrived to an empty queue with the server at his optimal resting location. If the server is currently serving a customer who arrived at radius r_i , then the

server's radial location is between $r_i - R t_i^{\max}$ and $r_i - R t_i^{\min}$ and as a rough estimate of the range of the server's radial location, we take a weighted average of these ranges:

$$r_{\min} = \sum_{i=1}^K \frac{\lambda_i}{\lambda} (r_i - R t_i^{\max}), \quad (\text{A.27})$$

$$r_{\max} = \sum_{i=1}^K \frac{\lambda_i}{\lambda} (r_i - R t_i^{\min}). \quad (\text{A.28})$$

If the server is located at generic location r_s at the time of customer arrival, the customer becomes uncatchable when he reaches radius $(r_s)/(\alpha)$. Our assumption that the server's location is $U[r_{\min}, r_{\max}]$ implies that T_{1i} , which is the time for a customer arriving to r_i to become uncatchable, is $U[r_{li}, r_{ui}]$, where

$$r_{li} = \frac{1}{R} \left(r_i - \frac{r_{\max}}{\alpha} \right)^+, \quad r_{ui} = \frac{1}{R} \left(r_i - \frac{r_{\min}}{\alpha} \right)^+. \quad (\text{A.29})$$

The renegeing probability p_r , where $p_n = 1 - p_r$ was used to compute the service rate μ , is given by $\sum_{i=1}^K (\lambda_i)/(\lambda) P(T_{1i} < T_2)$, which by (40) in [25] equals

$$p_r = \sum_{i=1}^K \frac{\lambda_i}{\lambda} \left(\frac{e^{-\mu_r r_{li}} - e^{-\mu_r r_{ui}}}{\mu_r (r_{ui} - r_{li})} \right). \quad (\text{A.30})$$

Equating $(\gamma)/(\gamma + \mu)$ to p_r in (A.30) yields our renegeing parameter,

$$\gamma = \frac{p_r \mu}{1 - p_r}. \quad (\text{A.31})$$

Finally, to compute the expected damage inflicted by a terrorist, we note that the terrorist makes it to the target undetected with probability $1 - \sum_{i=1}^K (1-d)^{i-1} d$. Among detected terrorists, a fraction p_2 are detected when there are already two customers in the system and a fraction $p_1((\gamma)/(\gamma + \mu))$ renege after being detected when there is one other customer in the system; in both cases, the terrorist makes it to the target and causes damage b . Among detected terrorists, a fraction p_0 arrive to an empty system, and if they are detected at radius r_i then they are caught on average at radius $r_i - d_i^e$, where d_i^e is given in (A.17), at which point they successfully detonate the bomb with probability q . Similarly, among detected terrorists, a fraction $p_1(\mu/\gamma + \mu)$ arrive to find one other customer in the system but are eventually caught, at which point they successfully detonate the bomb with probability q .

It remains to compute the mean radial location at which a non-renegeing terrorist is caught. The mean amount of time a non-renegeing terrorist arriving to radius r_i travels before the server begins chasing him is $E[T_2 | T_2 < T_{1i}^i]$. As in [25], we replace T_{1i}^i by its mean, $(r_{li} + r_{ui})/(2)$, in this conditional expectation, and assume T_2 is normal with mean m_r in (A.26) and variance σ_r^2 . By (42) in [25], the variance of the residual service time for a customer who arrived at radial location r_i is $(1/3)(t_i^e + m_s)^2 +$

$\sigma_s^2 + (1/12)(t_i^{\max} - t_i^{\min})^2 - m_{r,i}^2$. Therefore the variance is given by

$$\sigma_r^2 = \sum_{i=1}^K \frac{\lambda_i^2}{\lambda^2} \left(\frac{1}{3} (t_i^e + m_s)^2 + \sigma_s^2 + \frac{1}{12} (t_i^{\max} - t_i^{\min})^2 - m_{r,i}^2 \right). \quad (\text{A.32})$$

By (44) in [25], the non-renegeing terrorist's mean radial location at the time the server starts chasing him is

$$\tilde{r}_i = r_i - R \left(m_r - \frac{\sigma_r \phi \left(\frac{r_i + r_{ui} - m_r}{\sigma_r} \right)}{\Phi \left(\frac{r_i + r_{ui} - m_r}{\sigma_r} \right)} \right). \quad (\text{A.33})$$

We assume that the server is located at the mean radial location after catching a nuisance customer from his optimal resting location, which by (A.17) is

$$\tilde{r}_s = \sum_{i=1}^K \frac{\lambda_i}{\lambda} (r_i - d_i^e). \quad (\text{A.34})$$

By (A.17), the mean distance traveled by the non-renegeing terrorist during the chase is

$$d_i^b = \begin{cases} \tilde{r}_i - \frac{M}{\pi} (\alpha \tilde{r}_i - \tilde{r}_s) \ln \left(1 + \frac{\pi}{(\alpha-1)M} \right) & \text{if } \tilde{r}_i \leq \tilde{r}_s; \\ \tilde{r}_i + \frac{\alpha M (\tilde{r}_i - \tilde{r}_s)^2 (\alpha+2)}{2\pi \tilde{r}_s (\alpha+1)} - \frac{\alpha M \tilde{r}_i (\tilde{r}_i - \tilde{r}_s)}{\pi \tilde{r}_s} & \text{if } \tilde{r}_s < \tilde{r}_i < \tilde{r}_s \left(1 + \frac{\pi}{\alpha M} \right) \\ -\frac{M}{\pi} (\alpha \tilde{r}_i - \tilde{r}_s) \ln \left(\frac{\alpha-1+\pi/M}{\tilde{r}_i-1} \right) & \text{if } \tilde{r}_i \geq \tilde{r}_s \left(1 + \frac{\pi}{\alpha M} \right), \\ \frac{\tilde{r}_i - \tilde{r}_s (1 - \frac{\pi}{2M})}{\alpha+1} & \end{cases} \quad (\text{A.35})$$

and the mean radius at which he is caught is $\tilde{r}_i - d_i^b$. As noted earlier, we use (A.35) and set $t_i^b = (d_i^b)/(R)$ in computing the service rate μ of the $M/M/1/2$ queue.

Taken together, the expected damage, denoted by $E[D]$, is

$$E[D] = b + \sum_{i=1}^K (1-d)^{i-1} d \left[p_2 b - b + qp_0 [d_1 - d_2 (r_i - d_i^e)] + p_1 \left(\frac{\gamma b}{\gamma + \mu} + \frac{\mu q [d_1 - d_2 (\tilde{r}_i - d_i^b)]}{\gamma + \mu} \right) \right]. \quad (\text{A.36})$$

The accuracy of (A.36) is assessed in Fig. 4 and found to be very accurate for low-to-moderate utilization, which is the practical regime.

APPENDIX IV

THE OPTIMIZATION PROBLEM

The government's optimization problem is:

$$\min_{K, \bar{n}, M} E[D] \quad (\text{A.37})$$

subject to

$$d = P(X_1 > \bar{n}), \quad (\text{A.38})$$

$$f = P(X_2 > \bar{n}), \quad (\text{A.39})$$

$$k^* = \begin{cases} 0 & \text{if } d_1 - d_2 N \leq \sum_{i=N-K+1}^N \frac{q(d_1 - d_2 i)}{(N-i+1)(N-i+2)} + \frac{d_1}{K+1}; \\ N & \text{if } d_1 - d_2 N > \sum_{i=N-K+1}^N \frac{q(d_1 - d_2 i)}{(N-i+1)(N-i+2)} + \frac{d_1}{K+1}, \end{cases} \quad (\text{A.40})$$

$$B \geq c_Y Y + c_M M. \quad (\text{A.41})$$

ACKNOWLEDGMENT

L. M. Wein would like to thank T. Edmunds and R. Wheeler for helpful conversations.

REFERENCES

- [1] M. Bunn, A. Wier, and J. P. Holdren, *Controlling Nuclear Warheads and Materials: A Report Card and Action Plan*. Cambridge, MA: Harvard Univ. Press, 2003, Project on Managing the Atom, John F. Kennedy School of Government.
- [2] S. E. Flynn, "Beyond border control," *Foreign Affairs*, vol. 79, pp. 57-65, 2000.
- [3] R. M. Stana, Summary of Challenges Faced in Targeting Ongoing Cargo Containers for Inspection, U.S. General Accounting Office Rep. GAO-04-557T, Mar. 31, 2004.
- [4] U.S. Department of Homeland Security. Fact sheet: Domestic Nuclear Detection Office [Online]. Available: <http://www.dhs.gov/dhspublic/display?theme=43&content=4474&print=true>, Sep. 23, 2005
- [5] R. Gibbons, *Game Theory for Applied Economists*. Princeton, NJ: Princeton Univ. Press, 1992.
- [6] F. Pan, W. S. Charlton, and D. P. Morton, "A Stochastic Program for Interdicting Smuggled Nuclear Material," in *Network Interdiction and Stochastic Programming*, D. L. L. Ed. Norwell, MA: Kluwer, 2003, ch. Chapter 1.
- [7] L. Wein and M. Baveja, "Using fingerprint image quality to improve the identification performance of the U.S. visitor and immigrant status indicator technology program," *Proc. Nat. Acad. Sci. USA*, vol. 102, pp. 7772-7775, 2005.
- [8] L. M. Wein, A. H. Wilkins, M. Baveja, and S. E. Flynn, "Preventing the importation of illicit nuclear materials in shipping containers," *Risk Anal.*, vol. 26, pp. 1377-1393, 2006.
- [9] S. Fetter, V. A. Frolov, M. Miller, R. Mozley, O. F. Prilutsky, S. N. Rodionov, and R. Z. Sagdeev, "Detecting nuclear warheads," *Sci. Global Security*, vol. 1, pp. 225-302, 1990.
- [10] P. Beck, Illicit Trafficking Radiation Detection Assessment Program Paper OEFZS-G-0005, Austrian Research Centers, Seibersdorf (2000).
- [11] I. Weinsall, New York City Bridge Traffic Volumes 2001, Nov. 2002, Department of Transportation, New York City.
- [12] D. E. Archer, Adaptable Radiation Area Monitor (ARAM) (in Lawrence Livermore National Laboratory Technical Rep.) Livermore, CA, UCRL-JRNL-209101, 2005.
- [13] B. A. Reaves and M. J. Hickman, Law Enforcement Management and Administrative Statistics, 2000: Data for Individual State and Local Agencies With 100 or More Officers. Bureau of Justice Statistics, Office of Justice Programs, U.S. Department of Justice. Washington, D.C., 2004.

- [14] D. P. Bertsekas, *Dynamic Programming and Stochastic Control*. New York: Academic, 1976.
- [15] D. Gross and C. M. Harris, *Fundamentals of Queueing Theory*, 2nd ed. New York: Wiley, 1985.
- [16] B. Rooney, 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, Jun. 21, 2005.
- [17] E. Lipton, "U.S. to spend \$1.2 billion on detecting radiation," *NYTimes*, p. A10, Jul. 15, 2007.
- [18] E. W. Pontes and A. Ferreira, "Using cumulants and spectra to model nuclear radiation detectors," *IEEE Trans. Nucl. Sci.*, vol. 53, no. 3, pp. 1292–1298, Jun. 2006.
- [19] M. S. Gordon, P. Goldhagen, K. P. Rodbell, T. H. Zabel, H. H. K. Tang, J. M. Clem, and P. Bailey, "Measurement of the flux and energy spectrum of cosmic-ray induced neutrons on the ground," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 6, pp. 3427–3434, Dec. 2004.
- [20] TSA Systems, Ltd. Portal monitors, VM-250A/VM-250AGN [Online]. Available: http://www.tsasystems.com/products/portal_VM-250A.html, Aug. 30, 2007.
- [21] R. T. Kouzes, Pacific Northwest National Laboratory. Richland, WA, Jan. 3, 2007, private communication.
- [22] E. Normand and T. J. Baker, "Altitude and latitude variation in avionics SEU and atmospheric neutron flux," *IEEE Trans. Nucl. Sci.*, vol. 40, no. 6, pp. 1484–1490, Dec. 1993.
- [23] P. Goldhagen, "Cosmic-ray neutrons on the ground and in the atmosphere," *MRS Bulletin*, vol. 28, pp. 131–135, 2003.
- [24] W. Hage and D. M. Cirafelli, "Correlation analysis with neutron count distributions in randomly or signal triggered time intervals for assay of special fissile materials," *Nucl. Sci. Eng.*, vol. 89, pp. 159–176, 1985.
- [25] M. P. Atkinson, Z. Cao, and L. M. Wein, "Modeling and analysis of radiation sensor arrays around cities," 2007 [Online]. Available: <http://faculty-gsb.stanford.edu/wein/personal/radiation-wall.pdf>
- [26] S. Glasstone and P. J. Dolan, *The Effects of Nuclear Weapons 1977*, U.S. Dept. of Defense, Energy Research and Development Administration.
- [27] J. Berger, *Statistical Decision Theory and Bayesian Analysis*. New York: Springer-Verlag, 1985.
- [28] D. J. Bertsimas and G. van Ryzin, "Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles," *Oper. Res.*, vol. 41, pp. 60–76, 1993.
- [29] G. Grimmett, *Percolation*, 2nd ed. Berlin, Germany: Springer-Verlag, 1999.
- [30] D. Huizenga, 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, Jun. 21, 2005.
- [31] E. H. Kaplan and M. Kress, "Operational effectiveness of suicide-bomber-detector schemes: A best-case analysis," *Proc. Nat. Acad. Sci. USA*, vol. 102, pp. 10399–10404, 2005.
- [32] American Association for the Advancement of Science, DHS Receives Modest R&D Boost in Final 2006 Budget [Online]. Available: <http://www.aaas.org/spp/rd/dhs06c.pdf>, Nov. 3, 2005.