

Probabilistic Attack Reconstruction and Resource Estimation in “Reload” Scenarios

Paul T. Boggs
Sandia National Laboratories
Livermore, CA 94551
Email: ptboggs@sandia.gov

David M. Gay
Sandia National Laboratories
Albuquerque, NM 87185
Email: dmgay@sandia.gov

Jaideep Ray
Sandia National Laboratories
Livermore, CA 94551
Email: jairay@sandia.gov

Abstract—Faced with a bioterrorist attack with an aerosolized pathogen preparation, an emergency manager will need to make medical resource allocations early in the aftermath, with large uncertainties in the estimates of the resource demand, and in ignorance of what subsequent attacks might occur at other sites. In this work we use earlier results on the reconstruction of bioattacks from partial observations of the patient stream to calculate a probability density function of the resource demand. Thereafter, we employ an optimization-under-uncertainty approach to perform resource allocations for both a single-site attack and a multiple-site attack, i.e., a reload scenario. We test the approach in situations where resources are scarce. Our approach can develop allocations profiles that have the potential to reduce the probability of an extremely adverse outcome in exchange for a more certain, but less adverse outcome. We also explore the effect of placing limits on daily allocations.

I. INTRODUCTION

In the event of an aerosolized anthrax attack, emergency managers will face the problem of allocating scarce resources to treat those infected. If the attack is detected on environmental sensors, managers may have only a very approximate idea of when the attack happened, where the release was, and what the typical dose received was; in the absence of detection (e.g., release at an uninstrumented location) the first intimation of an attack will be the confirmed diagnosis of the first few victims of the attack. Thus emergency managers will have to make allocation decisions in the face of substantial uncertainty. Further, the attack may be one of a series of time-staggered attacks at different locations (the “reload” scenario [1]). However, the time-series of people presenting themselves at hospitals requesting treatment for severe respiratory problems forms a concrete source of data. It will be some time before anthrax is confirmed. In the meantime, more patients will have come forward for treatment. In this paper, we develop a tool that uses a probabilistic model of the *attack parameters* (i.e., the time of release, the number of people infected and the average dose received) to compute an optimal strategy for allocating scarce resources. Note that this work will focus on medical personnel, equipment and other resources whose transportation poses a logistical burden. Medicines and other prophylactic interventions, whose transportation requirements are much simpler, are not considered here. The formulation is general and can accommodate both single-attack and reload scenarios, allowing for an equitable (in a certain sense) allocation among

several targets.

In a previous work, [2], the authors developed a technique that uses a time-series of diagnosed patients at the hospital, along with some other assumptions, to calculate a probability density function (PDF) of the attack parameters. With only a short time-series, the PDF is rather broad, reflecting the high uncertainty at the beginning of an attack. As more data are collected on subsequent days and the time-series lengthens, the PDF narrows, as expected. In [2] we illustrated the possible use of this PDF to construct optimal allocations of resources under significant uncertainty. The technique there was quite simple, but showed that reasonable models could be created and solved. This approach could thus be used to provide a powerful tool to emergency managers to help them with such a crisis, or as a planning tool to help assess the adequacy of resources and logistics for various attack situations.

In this paper, we extend the work in [2] in several interesting and important ways. First we show how to formulate the problem as a two-stage stochastic optimization problem with recourse. This is a more general formulation than in [2]. Then we consider the use of several different objective functions that might be used to measure the success of any allocation strategy. We also consider two additional complications and show that the resulting computational efforts are still quite reasonable. The first is the reuse of resources in which we estimate the number of people who will die in spite of treatment and the extent to which we can apply those resources to new patients. The second is the possibility of a subsequent attack and the need to spread the available resources “equitably” over all of the targets.

It should be noted that certain aspects of the resource allocation problem are classical operations research problems that have been studied for years. Issues such as multiple supply depots, differing costs for shipment to different cities, time delays in shipping, etc., have all been thoroughly studied and can easily be incorporated into any of our models. To keep things simple here, we ignore most of these and simply consider what the optimal allocation schedule should be regardless of these other considerations.

It should also be clear that there is no one “right” way to formulate the objective function for this problem. We will consider some possible alternatives and comment on the meaning and implications that can be attached to each. Similarly,

as will become clear in the formulation of the constraints for the multi-target (reload) scenario, there must be some specification of how the resources are to be shared. We propose a way to do this, but make no claim that other constraints, e.g., social constraints, may not be more appropriate in other circumstances.

The paper is organized as follows. We begin by summarizing in Section II some previous related work. Then, in Section III, we provide the basic derivation of the stochastic optimization problem with recourse. This is done for the single target, or single city, case. This allows us to describe the problem with simpler notation, but affords the opportunity to introduce the relevant constraints and to discuss several objective functions. In Section IV we give some numerical results for the single city case that illustrate some of the differences. In Section V we discuss the extension to several cities and show some numerical results illustrating this case. Finally, in Section VI we give some conclusions and opportunities for future research.

II. PREVIOUS WORK

The problem of resource allocation during a bioterrorism event has generally been studied within a planning context. A set of scenarios is assumed and one studies the pros and cons of various intervention strategies. Wein et al [3] concluded that distributing prophylactics and setting up facilities for hospital care quickly (essentially, a surge in the medical treatment facilities at a location, with its attendant requirement for trained personnel, equipment and infrastructure) is the best way to mitigate casualties in the aftermath of an aerosolized anthrax attack. Other studies, e.g., [4], [5], have addressed optimal ways of executing the interventions. In these studies, the scale of the response is assumed to be known.

On the other hand, war-gaming exercises (e.g., the TOPOFF exercises [6]) involving bioterrorism have revealed that a key unknown in mounting a medical response to a release is the size/extent of the outbreak [7], [8]. While underestimating the scale of the required response can lead to disaster ([7], [8]), an overestimation can cause needless delay (and consequently, casualties) and disruption. Within the context of a reload scenario, it can also lead to a premature expenditure of scarce resources, leaving one unprepared to respond to a subsequent attack. To this end, there have been studies to estimate the size of an outbreak, given partial observations. In [9], using a dose-independent incubation period model (i.e., low-dose exposure), the authors used a Bayesian approach to infer the time and size of an anthrax outbreak. In [10], the authors, in conjunction with a simple plume model for aerosolized anthrax release, developed a Bayesian method to infer the location and time of a release, primarily to prime a syndromic surveillance system for early warning [11]. In [12], the authors developed a spatiotemporal approach to estimate the size and location of an aerosolized anthrax release, and quantified the uncertainty. As few as 10 diagnosed patients were sufficient to infer many of the attack characteristics to a specificity sufficient for mounting a response.

In [2], the current authors addressed the problem of estimating the attack parameters from a short time-series of diagnosed patients. A purely temporal approach was used to formulate a statistical inverse problem, which was solved using a Bayesian technique. The attack parameters were inferred as PDFs. These PDFs could be developed with a time-series three days long; with more data, the PDFs narrowed, indicating a reduction in the uncertainty in the inferences. The model was applied to the Sverdlovsk anthrax outbreak and was able to identify the time of the release with 3 days of data, and came close to estimating its size with a time-series 9 days long. The outbreak lasted 42 days. We also explored the possibility of using the inferred attack parameters to plan a medical response, i.e., to estimate the quantity of medical resources required. Using the joint PDF of the attack parameters, we developed an ensemble of possible epidemic realizations and allocated resources to minimize casualties, subject to the constraint that only a certain (insufficient) quantity of resources were available. These allocations were deemed “optimal” allocations and each realization of the epidemic was called a “scenario”. Thereafter, using a quadratic programming approach, we developed a “least-regret” allocation that would tend to minimize both resource wastage and casualties, given the uncertainty in how the outbreak would actually develop temporally. Regularization (in the optimization problem) was introduced as a penalty on the amount of resource used (i.e., we penalized large allocations as a way to hedge). It was observed that the least-regret allocation was very different from the “naive” allocation (an average over the optimal allocations for the members of the ensemble of epidemic realizations). The PDF of casualties arising out of the “naive” allocation was seen to be “long-tailed” while those arising from the least-regret allocation was peaked. Thus the least-regret allocation traded away a small risk of very high casualties for a smaller, but more certain level of deaths.

In [2], the rationale behind including a study of resource allocation lay in demonstrating how the characterization of the attack could be used profitably. It considered neither the reuse of resources (once an anthrax-infected patient had died), nor the inclusion of recourse in the models, i.e., consideration of corrections that can be made after new information has become available. It also did not consider the question of “reload”, i.e., time-staggered attacks, where allocation of resources have to be performed between two (or more) resource consumers, whose demands (and the degree of uncertainty in them) are different. We examine these more resource-allocation centric issues in this paper.

Before presenting our formulation of the problem, we remark on the nature of the treatment for anthrax. It is assumed that patients are diagnosed with anthrax only during their prodromal phase and will receive a powerful antibiotic. If this is the only intervention received, then approximately 85% of these people will die (this number is taken from the Sverdlovsk outbreak, where 69 out of 80 infected people died [13], [14]). If medical resources are available, the patient is assumed to have access to intensive care / surgical equipment for tracheocentesis / pleural drainage and will face a mortality

probability of 45%. This number is taken from the 2001 anthrax attacks where 5/11 people infected with inhalational anthrax died [15], [16]. In the remainder of this paper, we refer to people receiving treatment as those who get these aggressive measures.

III. STOCHASTIC OPTIMIZATION

As noted above, the basic problem is one of making optimal allocation of resources under significant uncertainty. In the previous section, we showed how to capture this uncertainty in a PDF that can be used to predict the number of patients who arrive at the hospital requesting treatment. Specifically, we construct a number of such scenarios that are consistent with the data we currently have.

The basic strategy is to use the data that we have to make a decision of how much resource to allocate today. We also estimate the quantity of resources that might be required in the future (i.e., a resource allocation profile) to plan future logistical requirements. When new information is obtained the next day, a new estimation of the attack parameters (and the resource demands) is performed. The scenarios are recalculated in light of the new data, the available resources are decreased by the amount allocated today, and the allocation recalculated. Thus we have one decision variable: the amount we allocate today based on the information at hand. Let that variable be a .

We adopt the following assumptions and notations: Let

- K be the number of scenarios;
- T be the number of days that we consider, i.e., the planning horizon of the epidemic;
- $N_{k,j}$ be the number of people requiring treatment who arrive on day j in scenario k ;
- $r_{k,j}$ be the allocation made on day j in scenario k ;
- $D_{k,j}$ be the number of people who die on day j in scenario k ;
- D_k be the total number of people who die in scenario k ;
- R be the total number of resource units available for the attack, where we assume for simplicity that one unit of resource treats one patient;
- $s_{k,j}$ be the number of resource units available on day j of scenario k ;
- t_j be the fraction of people, arriving on day j , who will die having been treated;
- u_j be the fraction of people, arriving on day j , who will die not having been treated.

We assume that those arriving later in the attack will be more likely to be successfully treated. This is motivated by the fact that longer incubations generally indicate a lower dose exposure (or a robust constitution). Thus we assume that $t_{j+1} > t_j$ and $u_{j+1} > u_j$. In practice, we make this difference ($t_{j+1} - t_j = u_{j+1} - u_j = \epsilon = 10^{-6}$) small and it merely serves as a mathematical stratagem to remove multiple solutions.

To construct the optimization problem, we need to specify the objective function. As a first cut, let us assume that we want to minimize some function of the sum of the number of

deaths in each scenario, i.e., we seek to minimize

$$\sum_{k=1}^K M(D_k)$$

where $M(D_k)$ is some measure of D_k . We could consider various measures, but clearly one could take M to be simply the expected number of deaths. The optimization problem is then

$$\min_a \sum_{k=1}^K M(D_k),$$

where we have to specify constraints on resources and on how to compute D_k . The resource constraint is, clearly,

$$0 \leq a \leq R.$$

Given the treatment assumptions described above, we can easily compute

$$D_{k,j} = r_{k,j}t_j + (N_{k,j} - r_{k,j})u_j. \quad (1)$$

For day 1, we tentatively substitute a for $r_{k,1}$.

The allocations $r_{k,j}$, $j > 1$ can be chosen to be the optimal allocations for scenario k , given that allocation a was made in day 1. These allocations will be constrained as follows:

$$\begin{aligned} r_{k,j} &\geq 0 \\ \sum_{j=2}^T r_{k,j} &\leq R - a \text{ for each } k. \end{aligned} \quad (2)$$

Although it is possible to iteratively solve problems for each scenario separately, it is more efficient to make the collection $r_{k,j}$ variables in the optimization problem and solve one large problem rather than K smaller problems for each trial value of a .

Before we pose the final version of the initial problem, we must address an important situation. It is possible that in some scenarios, an allocation a or $r_{k,j}$ will be greater than the number of people who arrive, i.e., $a > N_{k,1}$, in which case, the value of $D_{k,j}$ from above will not correctly calculate the number of deaths. To handle this situation, instead of substituting a for all $r_{k,1}$, we retain separate $r_{k,1}$ variables and impose the constraint $r_{k,1} \leq a$ for all k . We also impose the simple bound constraints $r_{k,j} \leq N_{k,j}$ for all (k, j) , and we change (2) to

$$\sum_{j=2}^T r_{k,j} \leq R - r_{k,1}.$$

Another important concern is that without further constraints, the optimal choice of a may be to allocate all possible resources on the first day, which seems unlikely to be the best policy. One way to address this issue is to make tentative allocations for all days in the planning horizon, i.e., to introduce decision variables $a_i \geq 0$ for $1 \leq i \leq T$ and to restrict each scenario's allocations by $r_{k,j} \leq a_j$, with $a_1 = a$ and

$$\sum_{j=1}^T a_j \leq R.$$

In other words, we decide, *a priori*, that the daily allocation cannot exceed a certain level. Obviously the level chosen has a significant impact on the quality of the allocation calculated. This is studied further below.

Of course, the purpose of the exercise is still to choose the first day's allocation $a = a_1$. Another possibility is to penalize over-allocation of resources, in keeping with some standard approaches. To do this, we introduce a penalty term in the objective function of the form

$$\rho \cdot (a - r_{k,1})_+,$$

where $x_+ = x$ if $x > 0$ and 0 otherwise, and ρ is a constant chosen to appropriately balance the costs, i.e., penalize wastage / overallocation of resources. We choose this form for our studies here.

Policy makers may further wish to limit daily allocations to specified fractions of the available resources, say $a_j \leq \sigma_j R$. For simplicity, below we use a common value $\sigma_j = \sigma \in (0, 1]$ for all j (with $\sigma = 1$ imposing no further restriction).

The final topic we consider here is the reuse of resources. As noted, a high percentage of patients being treated will die anyway and they will die at a nonuniform rate. Some, in fact, will die quite early and their resources can be used on incoming patients. Data for estimating the rates are not readily available, but reasonable approximations can be made. Based on typical treatment progressions, the longer one survives, the more likely complete recovery becomes. Thus the percentage of people who die after n days of treatment should increase rapidly for a few days and then gradually decrease. As a first cut, we assumed a ten-day period and used a simple function, $f_n = f_n^0 / \sum_{k=1}^{10} f_k^0$ with $f_n^0 = n / (1 + \exp(n/2))$ to estimate these rates (with $f_n = 0$ for $n > 10$). This is in the form of the expected percentage of people being treated who will die n days after treatment has begun. It is straightforward to then estimate the number of resources that will be available on any given day, as in (5) below. Along the same lines, as noted above, there will be some scenarios for which allocations will exceed demand and the extra resources will likewise be available for incoming patients. The number of people who will die is still given by (1).

Our optimization problem is shown in Figure 1. Some remarks about it are in order. It is a two-stage stochastic optimization problem with recourse. The first stage is today and the second stage is days 2– T . Each scenario takes recourse on the basis of today's allocation and does the best that it can after that. Constraints (3), (4), and (5) together imply that each scenario consumes at most R resources.

One could, in principle, construct a multi-stage problem by dividing the days 2– T into two or more stages. Suppose, for example, that the second stage is days 2–4. Then one could trace each of the K scenarios through day 4. At that point, one assumes that, for each k , the data $N_{k,j}$, $j = 1, \dots, 4$, are "true", constructs a PDF based on this data and samples that to obtain K new time series for each k . Although this can be easily continued, it is clear that the number of possible paths through the attack grows rapidly. In this paper, we confine

$$\begin{aligned} \min_{a, r_{k,j}, s_{k,j}} \quad & \frac{1}{K} \sum_{k=1}^K \{M(D_k) + \rho(a - r_{k,1})_+\} \\ \text{subject to:} \quad & 0 \leq a \leq \sigma R \\ & 0 \leq r_{k,j} \leq \min(N_{k,j}, \sigma R) \\ & r_{k,1} \leq a \\ & r_{k,j} \leq s_{k,j} \quad (3) \\ & s_{k,1} = R \quad (4) \\ & s_{k,j} = s_{k,j-1} - r_{k,j-1} \quad (5) \\ & \quad \quad \quad + \sum_{n=1}^{j-1} f_n t_{j-n} r_{k,j-n} \\ D_{k,j} = \quad & t_j r_{k,j} + (N_{k,j} - r_{k,j}) u_j \\ D_k = \quad & \sum_{j=1}^T D_{k,j}. \end{aligned}$$

Fig. 1. Multi-scenario resource allocation problem.

ourselves to just the two stages, but the extension to more stages is theoretically possible.

We have not yet specified M in the objective function, but we note that the constraints are all linear. Thus, if M is a linear function, we have a classical linear programming problem for which there are many excellent algorithms available. If we take M to be the identity operator, i.e.,

$$M(D_k) = D_k = \sum_{j=1}^T D_{k,j} \quad (6)$$

then we are simply computing the expected number of deaths in each scenario, and

$$\bar{D} = \frac{1}{K} \sum_{k=1}^K D_k \quad (7)$$

is the expected number of deaths over all the scenarios. This has an obvious appeal; results using this choice of M are reported in the next section. A potential problem with this is that scenarios with a large number of infected people could dominate the decisions. Recall our assumption that people arriving later are better candidates for treatment; in a scenario with a large number of people, the algorithm would delay the allocation of resources much more so than for a scenario with a much smaller number of infected people. It could be argued that the sampling procedure should properly account for this, but one could also divide D_k by the total number of people infected in scenario k . This downplays the influence of the larger cases, while keeping the problem linear.

A different approach, related to the work in [2], is to compute the optimal number of deaths for each scenario in K separate problems. Call the results D_k^* . Then one could obtain an allocation that stays as close as possible to all of these in some sense. A natural way to do this is to minimize the

variance between the vector D_k^* and the vector D_k resulting from any other allocation. In particular, one could use

$$\sum_{k=1}^K (D_k - D_k^*)^2$$

as the objective function. In [2] we referred to this as the least-regret formulation with the interpretation that the allocation made today is the one that we will least regret in the future since it does reasonably well for all scenarios. As above, we could scale each of the terms by the total number of arrivals in that scenario. Since this is a quadratic function, the optimization is now a convex quadratic programming problem; again, good algorithms exist. The computation of each D_k^* is a small linear programming problem that is solved quickly.

We now present some numerical results illustrating some of the issues raised here.

IV. RESULTS FOR ONE CITY

To explore the approach, we first generated a test case involving an anthrax attack on a single city. This is described in [2]. Briefly, an aerosolized anthrax release is simulated over a domain with spatially variable population density. Per this distribution and an atmospheric dispersion model, 22,384 individuals are infected with a range of doses, with an average dose of 1470 spores. People develop symptoms over time; the time series for the first 10 days is $\{3, 123, 719, 2046, 2202, 2194, 2058, 1918, 1656\}$. This time series was used to draw 100 samples from the joint PDF of the attack parameters using a single-component random-walk Markov Chain Monte Carlo (MCMC) sampler. Note that these samples were drawn after the MCMC sampler had “burnt-in” and had “converged” per the `mcgibbsit` package in R (Chapters 7 and 8 in [17]; also see [18]). For each attack parameter sample, 10 epidemic realizations were calculated (the forward model is stochastic), resulting in a set of 1000 epidemic realizations (or scenarios). Such ensembles, generated from the first 5 days of data in the time-series above, are plotted as the gray region in Figure 2. Note that we measure time from the day that the first person was diagnosed with anthrax (rather than the time of attack/infection). The distribution developed with data collected through Day 7 is much narrower than that through Day 3, confirming that the addition of 4 extra days of data significantly reduces the uncertainty. This has not been plotted here.

The model was implemented in AMPL [19]–[21] and used the CPLEX 11 [22] simplex method to solve the problems.

We ran many tests based on the model described above. We fixed our available resources such that they could treat 10,000 patients (out of the 22,384 infected), i.e., they are scarce. Our first observation is that the form of the function M does not make much of a difference in the results. Thus all of the results we show here were calculated using (6) to minimize the expected deaths (7). Our second observation is that the penalty parameter, ρ , should be taken to be a small value to ensure its desired effect. After some tests with several values

of ρ , summarized in Table I, we chose use $\rho = 0.001$ for all of the results reported here.

ρ	a	\bar{D}
0.	10000	9358.0
0.0001	2384	9358.0
0.001	2364	9358.0
0.01	2317	9358.2
0.1	2261	9360.2

TABLE I
EFFECT OF ρ ON a AND EXPECTED DEATHS (7).

In Figure 2 we plot the allocations, given a resource demand drawn from 5 days of observations in the time series. The gray region denotes the ensemble of scenarios. The time-series values used for the inference are plotted with triangles; the future observations in the time-series are plotted with diamonds. Allocations were calculated for $\sigma = 0.04$ and 0.1. Clearly σ makes a significant difference. Recall that there are two possible reasons for imposing a constraint on the amount of resource that can be shipped on a given day: the first is that this may simply be a logistical constraint; the second is that the emergency manager may want to conserve resources as a hedge against a subsequent attack. Observe that our formulation only computes the allocation for Day 6; to give managers an idea of allocations that might be appropriate on subsequent days, we obtain tentative allocations for days 7– T by averaging the allocations for each day over all of the scenarios. (Subsequently arriving data should influence the actual allocations for later days.) As is evident in Figure 2, the severe restriction imposed by $\sigma = .04$ implies that many fewer resources can be allocated than for the lighter restriction of $\sigma = .10$. Thus there is a commensurate increase in the number of deaths with $\sigma = .04$, as we show in Figure 3. Here we plot the PDF of excess casualties (over the optimal/minimal level that we would have achieved had we perfect knowledge of the epidemic) for the two values of σ . As might be expected, the effect of σ (i.e., the placing of a ceiling on how much can be shipped on a given day) is felt mainly in those scenarios that project a large number of infected people turning symptomatic. We also see that increasing σ narrows the PDF (we reduce the long-tail probability of an extremely adverse outcome) while raising the peak of the PDF and moving to lower values of excess casualties, i.e., increasing the probability of a less adverse outcome. Since the probability mass under the PDF is 1, this is tantamount to increasing the probability of a certain (acceptable) level of casualties while simultaneously trading it to reduce the probability of an extremely adverse outcome - a classic hedging / risk management operation. This is captured quantitatively in the change of shape of the PDF with σ .

For $\sigma \geq 0.20$ we obtain an allocation (not shown here) resulting in very few excess casualties in each scenario. These results show that the model can be used for assessing the effects of conserving resources in anticipation of a second attack or for planning purposes to see the need for a higher shipping capacity.




Fig. 2. Allocations for the attacked city, obtained from data collected over the first 5 days. The gray region denotes the evolution of all the scenarios considered. The net effect of σ is to reduce the allocation during the early days of the epidemic. $R = 10,000$. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds.

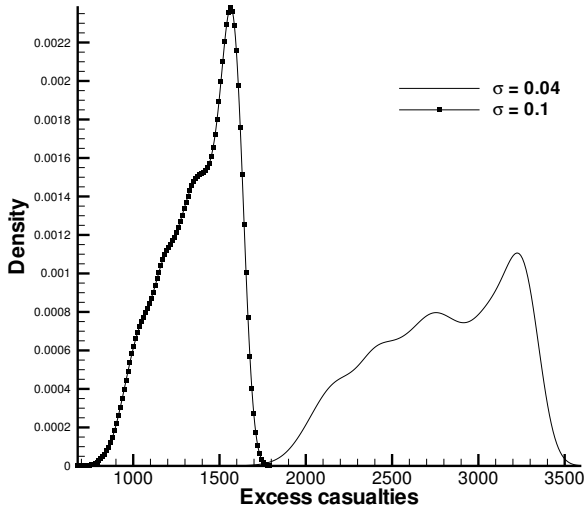


Fig. 3. PDFs of excess casualties for $\sigma = 0.04$ and 0.1 . $R = 10,000$. A tighter daily constraint on allocations ($\sigma = 0.04$) increases the probability of excess casualties. However, note that the PDFs have rather compact support.

One could also compare the PDFs of excess casualties if a “naive” approach to resource allocation was considered, e.g., given an R , one allocates on a scenario-by-scenario basis (leading to 1,000 allocations), then simply uses the mean of these allocations. Such a “naive” allocation results in a very long tail (see [2] for a comparison) and is not very competitive for hedging purposes vis-à-vis the more sophisticated techniques considered here and in [2]. For the

rest of this paper, the “naive” approach will be ignored.

V. THE RELOAD CASE

The main complication in dealing with an attack on two or more cities is in deciding how to allocate the resources among all of the cities. From the point of view of the model, it does not make any difference if a life is saved in the first city or the second. Thus, without further constraints, there is an inherent non-uniqueness in the solution of the problem, since the optimal number of deaths can be achieved in many ways, including the extreme one of sending all of the resources to one city and ignoring the other. In practice, it seems reasonable to assume that there will have to be some “social” or infrastructural constraint to ensure that all cities are treated fairly. We illustrate how this could be achieved below, but first we deal with another issue, namely that of whether or not to anticipate subsequent attacks.

As noted in Section III an emergency manager may wish to restrict the amount of resources that can be shipped on each day. This is done by imposing the constraints $r_{k,j} \leq \sigma R$. The manager could equally well choose σ to conserve some of the resources in case there is a subsequent attack, the “reload” case. If there is a subsequent attack, there is no way to say anything about it until there is some evidence in the form of people in the second city arriving at the local hospital in need of treatment. As is the case for the first city, a few days of data are required before any reasonable PDF can be computed and sampled.

Extending the basic model above to the case of several cities is straightforward. The major addition for the reload case is the social constraint. We illustrate the possibilities with a simple constraint that seeks to ensure that each city receives a proportional amount of the resources. A way to do this is to impose the constraints

$$D^i/A_i \leq (1 + \pi) \sum_{j \neq i} D^j / \sum_{j \neq i} A_j,$$

where D^i are the deaths in city i , A_i is the total number of patients in city i , and $\pi \in [0, 1]$. For the results reported here, we used $\pi = 0.1$, so that the relative resource allocations are within 10%.

We demonstrate this allocation approach on a simulated reload scenario. The first attack (on City A) is the same as in Section IV. However, on Day 3 of the first attack, City B records an anthrax diagnosis and it is verified that it too has been attacked. The time-series for City B is $\{0, 0, 1, 76, 711, 1765, 2720, 3099, 3186, 2896\}$ for the first 10 days. The attack on City B was simulated in the manner described in [2]. 29,861 people were infected, with an average dose of 2749 spores. The two attacked cities therefore have a resource demand of around 50,000 units. In the study below, we will assume that only 25,000 units are available.

The allocations are shown in Figures 4 and 5 for Day 6 of the attack, i.e., we have a time-series 5 days long for City A and 3 days long for City B. The gray region in Figures 4 and 5 show the ensemble of scenarios for the two attacks; as

expected, the ensemble for City B is far broader than City A, denoting a larger uncertainty arising from a smaller time-series of observations. The observed and unobserved evolution of the epidemic in the two cities is plotted using triangles and diamonds. The allocations developed with $\sigma = 0.04$ and 0.1 are plotted for Day 6 (and beyond) of the epidemic. Note that the allocation is only meant for Day 6. Both the plots demonstrate how allocations are curtailed as σ decreases, leading to extra casualties, especially for scenarios that project larger epidemics. Also note that the effect of σ is felt mostly during the peak of the epidemic; the allocations are similar towards the end. This is a consequence of our modeling decision to slightly favor later allocations.

In Figure 6 we plot the PDFs of excess casualties (over the minimum that we would achieved had we perfect knowledge of the attack and the epidemic). The excess casualties for Cities A and B, for $\sigma = 0.02, 0.04$ and 0.1 are totaled and plotted. Note that the $\sigma = 0.1$ case is not at all restrictive and one even has overallocation of resources (the “negative” casualties). This happens when two exceptionally small scenarios for City A and B are considered. Note that the σ value merely places a bound on daily allocation; the constraint that daily allocations must add up to the available resources is not violated. The hedging effect of σ seen in Section IV is also reproduced here, though with a few modifications. In all cases, we see a multimodal excess-casualty distribution. While $\sigma = 0.04$ does manage to translate the excess-casualty PDF to the left (vis-à-vis $\sigma = 0.02$), we see the width of its support is unchanged, i.e., the higher value of σ reduces the expected casualties (and consequently risk), but does not improve the hedge compared to $\sigma = 0.02$.

VI. CONCLUSION

We have developed a novel approach to the problem of allocation of scarce resources in the face of great uncertainty following an anthrax attack. Our approach involves the calculation of a PDF that captures the uncertainty in the nature of the attack.

We have shown that this PDF can be used to generate realistic scenarios regarding the evolution of the ensuing epidemic. These scenarios can be used within optimization models to provide reasonable resource allocation schedules. We show that these are likely to be effective over all of the scenarios generated. Our models contain constraints that add realism to them and show that they can be extended to address a wide variety of situations. They also contain a parameter σ that places a constraint on the maximum level of daily allocations. This serves as a proxy for an emergency manager’s risk appetite (e.g., if a second attack is expected and resources have to be conserved) as well as infrastructural limitations.

We have exercised the models in single- and two-attack cases. In the two-attack case, we have demonstrated how a degree of fairness in the allocations to the two attacked sites can be encoded into the optimization problem. We find that σ has a significant effect the optimality of the solution. However, this sub-optimality (alternatively, the risk of excess casualties)

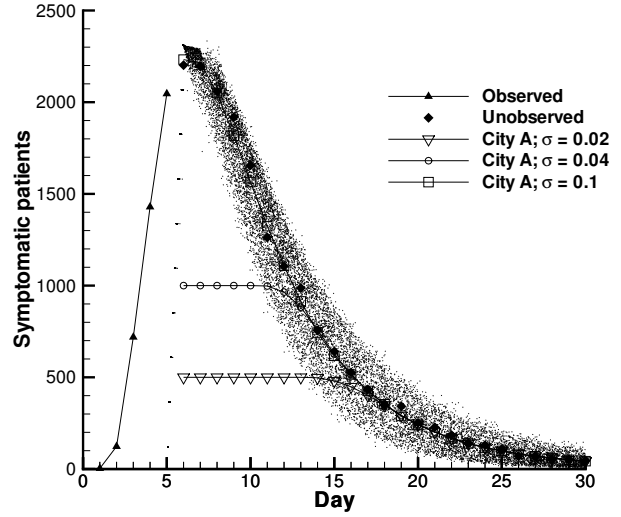


Fig. 4. Allocation under various values of σ for City A. The gray region denotes the evolution of all the scenarios considered. $R = 50,000$ (total for both cities). These allocations were drawn from data collected over 5 days; allocations are for Day 6. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds.

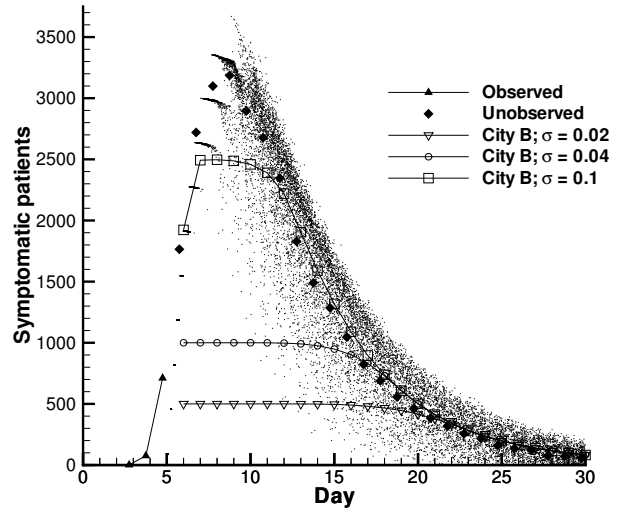


Fig. 5. Allocation under various values of σ for City B. $R = 50,000$ (total for both cities). These allocations were drawn from data collected over 3 days (Days 3, 4 and 5, the attack stagger being 2 days); allocations are for Day 6. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds. The gray region denotes the evolution of all the scenarios considered.

can be rigorously captured as a PDF; lower values of σ tend to lengthen the tail of the PDF creating a low, but non-zero possibility of extremely adverse outcomes. Thus, in a decision-theoretic setting, one could arrive at a recommended σ value based on a utility function, predicated on the PDF of casualties.

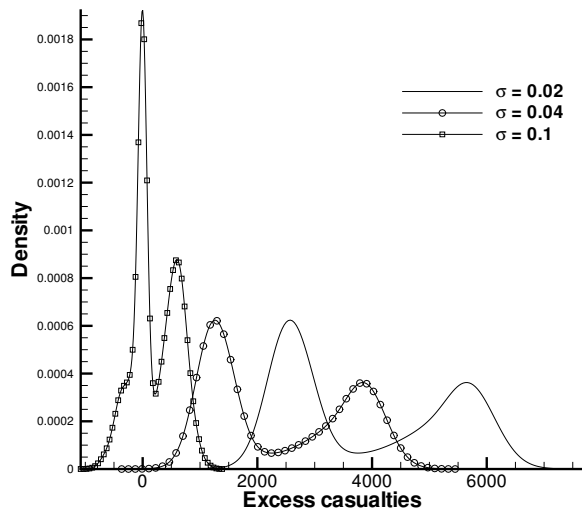


Fig. 6. PDF of excess casualties for $\sigma = 0.02, 0.04$ and 0.1 . $R = 50,000$. Note how $\sigma = 0.1$ results in certain “negative” casualties, i.e., excess resources in certain low casualty scenarios. Most of the excesses occur in the smaller attack (City A). The PDF corresponding to $\sigma = 0.02$ and 0.04 have a similar support widths but the higher value of σ reduced the expected value of casualties.

Our allocation models are general and could easily be extended to handle a complex network of supply depots and shipping strategies. We have shown how one may balance the allocations in a socially fair way in the multiple-attack case. The point of these extensions was not to be exhaustive, but to illustrate the fact that many constraints can be added to provide the emergency managers with a useful and powerful tool to help them make difficult decisions, or to plan over a wide variety of constraints.

Future work on this topic includes extensions that allow more complicated classes of resources with appropriate models of how these resources affect the number of survivors. More interesting is the consideration of a different type of disease. A simplifying feature of anthrax is that it is not contagious, so the scenarios that we generate are independent of any resource allocation strategy. In a contagious disease, however, resources allocated, along with other response strategies, e.g., quarantine of those infected, affect the progression of the epidemic. This coupling will be challenging to model, but we believe that the methods described here provide a foundation for addressing these cases.

ACKNOWLEDGMENT

The authors would like to thank Todd Plantenga for useful discussions on the formulation of the problem considered here. This work was supported by Sandia National Laboratories’ LDRD (Laboratory Directed Research and Development) program, with funds allocated under the “Enabling Predictive Simulation” Investment Area. Sandia is a multiprogram lab-

oratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy’s National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

REFERENCES

- [1] R. Danzig, *Catastrophic bioterrorism: What is to be done?* Bernal Press, PA, USA, 2003.
- [2] J. Ray, P. T. Boggs, D. M. Gay, Y. M. Marzouk, and H. N. Najm, “A Bayesian approach for estimating bioterror attacks from patient data,” *Statistics in Medicine*, 2009. Under review.
- [3] L. M. Wein, D. L. Craft, and E. H. Kaplan, “Emergency response to an anthrax attack,” *Proc. Natl. Acad. Sci. USA*, vol. 100, no. 7, pp. 4346–4351, 2003.
- [4] R. Brookmeyer, E. Johnson, and R. Bollinger, “Modeling the optimum duration of antibiotic prophylaxis in an anthrax attack,” *Proceedings of the National Academy of Science*, vol. 100, no. 17, pp. 10129–10132, 2003.
- [5] R. A. Fowler, G. D. Sanders, D. M. Bravata, B. Nouri, J. M. Gastwirth, D. Peterson, A. G. Broker, A. M. Garber, and D. Owens, “Cost-effectiveness of defending against bioterrorism: A comparison of vaccination and antibiotic prophylaxis against anthrax,” *Annals of Internal Medicine*, vol. 142, pp. 601–610, 2005.
- [6] D. “Transcript of background briefing with senior DHS officials on TOPOFF 3,” 2005. http://www.dhs.gov/xprepresp/training/gc_1179430526487.shtm, accessed March 2008.
- [7] T. O’Toole, M. Mair, and T. V. Inglesby, “Shining light on Dark Winter,” *Clinical Infectious Diseases*, vol. 34, pp. 972–983, 2002.
- [8] T. V. Inglesby, R. Grossman, and T. O’Toole, “A plague on your city: Observations from TOPOFF,” *Clinical Infectious Diseases*, vol. 32, pp. 436–445, 2001.
- [9] J. Walden and E. H. Kaplan, “Estimating time and size of bioterror attack,” *Emerging Infectious Diseases*, vol. 10, no. 7, pp. 1202–1205, 2004.
- [10] W. R. Hogan, G. Cooper, G. L. Wallstrom, and M. Wagner, “An Improved Bayesian Aerosol Release Detector,” tech. rep., The RODS Laboratory, 550 Cellomics Building, 100 Technology Drive, Pittsburgh, PA 15219, 2005.
- [11] W. Hogan, G. F. Cooper, G. Wallstrom, M. Wagner, and J. M. Depinay, “The Bayesian aerosol release detector: An algorithm for detecting and characterizing outbreaks caused by an atmospheric release of *Bacillus Anthracis*,” *Statistics in Medicine*, vol. 26, no. 29, pp. 5225–5252, 2007.
- [12] J. Legrand, N. Ferguson, S. Leach, I. Hall, S. Cauchemez, and J. Egan, “Estimating the location and timing of a covert anthrax attack,” *Public Library of Science - Computational Biology*, vol. 5, no. 1, p. e1000356, 2009.
- [13] M. Meselson, J. Guillemin, M. Hugh-Jones, A. Langmuir, I. Popova, A. Shelokov, and O. Yampolskaya, “The Sverdlovsk anthrax outbreak of 1979,” *Science*, vol. 266, pp. 1202–1208, 1994.
- [14] D. Wilkening, “Sverdlovsk revisited : Modeling human inhalational anthrax,” *Proceedings of the National Academy of Science*, vol. 103, pp. 7589–7594, May 2006.
- [15] J. A. Jernigan et al., “Bioterrorism-related inhalational anthrax: The first 10 cases reported in the United States,” *Emerging Infectious Diseases*, vol. 7, no. 6, pp. 933–944, 2001.
- [16] D. B. Jernigan et al., “Investigation of bioterrorism related anthrax, United States, 2001: Epidemiological findings,” *Emerging Infectious Diseases*, vol. 8, no. 10, pp. 1019–1028, 2002.
- [17] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *Markov Chain Monte Carlo in Practice*. Chapman and Hall, 1996.
- [18] “CRAN site for the mcgibbsit package.” <http://cran.r-project.org/web/packages/mcgibbsit/index.html>.
- [19] R. Fourer, D. M. Gay, and B. W. Kernighan, “A modeling language for mathematical programming,” *Management Science*, vol. 36, no. 5, pp. 519–554, 1990.
- [20] R. Fourer, D. M. Gay, and B. W. Kernighan, *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press/Brooks/Cole Publishing Co., second ed., 2003.
- [21] <http://www.ampl.com/>.
- [22] <http://www.ilog.com/products/cplex/>.