from *SIAM News*, Volume 35, Number 10, December 2002

"Hot Topics" Workshop Takes a Logistical Look at Biodefense

By Barry A. Cipra

Can mathematical modelers help the United States prepare for a bioterrorist attack? A "hot topics" workshop, sponsored by the Institute for Mathematics and Its Applications at the University of Minnesota and SIAM, held September 28 at the IMA, elaborated on the affirmative answer.

The workshop, organized by IMA director Doug Arnold, Mac Hyman of Los Alamos National Laboratory, and Edward Kaplan of Yale University, highlighted the interface between traditional epidemiological models of infectious diseases and operations research analyses of the logistical aspects of biodefense. The former delineates the likely course of events in a bioattack under various scenarios. The latter aims to identify optimal allocations of resources for a given set of objectives and constraints. Together, they can point toward strategies for dealing with bioterrorism—perhaps by persuading the evildoers, who have resource-allocation problems of their own, to direct their malevolent efforts elsewhere.

The biodefense workshop was held during the IMA's fall quarter on supply chain and logistics optimization. "It's a fascinating area," says Kaplan. "When you're talking about logistics and bioterrorism, very interesting things start to happen."

Epi-queuers

Traditional epidemiological models use differential equations to study the spread of disease in populations of susceptible, infected, and recovered (or dead) individuals. They help policy-makers judge the likely costs and benefits of various public health measures, such as needle exchanges for hepatitis and HIV, mosquito control for West Nile virus, or priority ordering for recipients of flu vaccines. However, the models usually assume an instantaneous implementation of whatever policy measures they include as parameters. This isn't normally a problem, either because the time scale of the epidemic is much longer than the ramp-up time of the policy or because preventive measures are (theoretically) in place before the outbreak of disease. In the case of flu, for example, the basic idea is to inoculate a sufficient percentage of the right people before the season starts, to prevent that year's anticipated strain from decimating the population.

But a bioterrorist attack takes place on a swifter time scale. Consequently, efforts to cope with the malicious release of a pathogen will compete with the underlying dynamics of the disease it produces. Moreover, the uncertainties of when, where, what kind, and even whether an attack will occur complicate the preparation problem. Models that compare the spread of disease through unvaccinated and, say, 50% vaccinated populations are inadequate, says Kaplan. "You instead have to say, 'How does the disease spread if you are simultaneously trying to vaccinate the population while it's spreading?' "

Enter logistics. An emergency response of any kind, be it to bioterror or hurricane Floyd, amounts to putting appropriate people and supplies in place to aid a distressed population. In the case of a bioattack, the response includes identifying sick people who need treatment (and possibly isolation) and doling out prophylactic measures to those at risk. The trick is to do this as quickly as humanly possible. "With the kinds of worries we face here, we're talking about very fast moving events, where a difference of a day or even, say, six hours, can lead to very different results," Kaplan says.

Queueing theory—picture long lines at a makeshift vaccination clinic—has a lot to say about what's humanly possible, as does supply chain analysis. Among the questions policymakers need to ponder: How many (and what kind of) responders do you need to mobilize in the event of an attack, and what should you have them do?

Calculus of Variolations

Take smallpox, for example. The federal government has ordered the stockpiling of smallpox vaccine sufficient to vaccinate the entire country. If the threat of such an attack were clear, the strategy for dealing with it would be simple and effective: Inoculate everyone pre-emptively. This is essentially how the disease was eradicated in the United States more than fifty years ago. (The last case of smallpox in the U.S. was reported in 1949. Routine vaccinations were discontinued in 1971.) But a "mass" inoculation program comes at a cost. It's estimated that for every million recipients of the vaccine, roughly ten will develop serious complications and one will die.

Given the current, extreme uncertainties about the availability of smallpox to terrorists and "rogue" states, policy-makers are loath to recommend preventive measures that would create casualties of their own. (When a natural outbreak of smallpox was still a clear and present danger, the risk was considered worth taking. Besides, back then, the vaccine was claiming only a few lives each year; a biodefense push to inoculate the entire nation today would kill scores in a matter of weeks. The media—not to mention the lawyers—would have a field day, especially since the adverse reactions to smallpox vaccine are visually horrifying.) Instead, current policy calls for vaccinations of the general population only if an attack occurs—only, that is, if a case of smallpox is confirmed.

Even then, there are two possible approaches. One is mass vaccination, in which the entire population is inoculated as quickly as possible. The other is "ring" or "traced" vaccination, in which resources are focused on tracking down the contacts in confirmed cases, the idea being to contain the outbreak by immunizing a ring around it, switching to mass vaccination only if the containment fails. (Actually, either strategy involves a screening step, because many people, including those with compromised immune systems or skin conditions like eczema, should not be inoculated because they are at high risk for complications from the vaccine. Screening is a potential bottleneck in any bioattack response.)

Ring vaccination proved effective in the endgame with smallpox in the 1970s (the last documented case occurred in Somalia in 1977), and the Centers for Disease Control initially proposed that it be used in the event of a bioattack. But an analysis by Kaplan and colleagues David Craft and Lawrence Wein of MIT, published in August in the *Proceedings of the National Academy of Sciences*, helped persuade policymakers that mass vaccination looks like a better strategy. Their model compares the two approaches, taking into account the inherent queueing delays of each. In their "base" case (1000 people initially infected in an urban population of 10 million), the number of fatalities is an order of magnitude less for mass vaccination than for the ring strategy, reducing the death toll from several thousand to several hundred. (The model also indicates that a pure ring strategy, with no switch to mass vaccination, is a terrible idea. In the base case, it leads to about 100,000 deaths. With no vaccine at all, smallpox, with its 30% mortality rate, would run amok in the model's population, wiping out three million people.)

The model consists of 17 differential equations, four of which are devoted to the flow of people through the vaccination queue. Using estimates from other public health studies, the modelers assume that 5000 vaccinators are working around the clock. With the mass strategy, each vaccinator can treat 200 people per day; the number drops to 50 per day with the ring approach, because the vaccinators spend three quarters of their time locating contacts. These numbers mean that the model's mass vaccination is complete in 10 days. The comparison with the ring strategy assumes that the switch to mass vaccination is made after four weeks of ring vaccination.

The precise numbers coughed up by the model are less important than the model's robustness. According to the PNAS paper, the superiority of mass vaccination holds over a range of parameter values. Ring vaccination is preferable only if the initial size of the outbreak is small (a dozen cases or so) and the basic reproductive ratio (the average number of people an infected person passes the disease along to early in the outbreak, which in the model's base case is taken as three) is low; neither of those measures is easily determined.

"Sensitivity analysis is important, because you want to be able to come to a policymaker with some robust conclusion," Kaplan says. "We're not claiming we know the numerical answer. We are trying to say we have a very good argument that policy A is better than policy B. And that's really what operations researchers try to do—they try to help people make better decisions." Besides, he says, "What are [policymakers] doing when they don't have these models? They're making decisions anyhow!"

Sporadic Numbers

Anthrax presents similar problems, but with important differences. Unlike smallpox, anthrax is not contagious, and it's curable if caught in time (but evenly deadlier if not). But a massive attack, say by an urban crop duster, could infect millions of people all at once. And anthrax is *known* to exist in weaponized form.

Detecting an anthrax attack and judging its extent are in themselves nontrivial problems. But the logistics of response makes them look easy. According to Wein (who has worked with Kaplan, both on the smallpox model and on a model for anthrax now being developed), the key bottleneck is likely to be hospital capacity and the number of emergency doctors. In other words, doling out antibiotics is no big deal, but treating hordes of seriously sick people will severely stress, and possibly overwhelm, the public health system—unless steps are taken to expand the number of trained responders and make them mobile. "Our security against bioattack rests largely on having a good response," Wein says.

Detecting an anthrax attack and judging its extent are in themselves nontrivial problems. But the logistics of response makes them look easy.

One of the lessons of last year's anthrax scare is that, while anthrax isn't contagious in people, it can be passed from letter to letter. At the IMA workshop, Glenn Webb, a mathematician at Vanderbilt University, described a model he and Martin Blaser of the New York University School of Medicine had devised to analyze mailborne transmission of the bug. Their model, published last May in PNAS, simulates the cross-contamination of mail as bug-laced letters pass through the postal system.

The model distills mail handling down to five stages, from pickup at point of origin (the corner mailbox) to local post office, through a regional center, to the recipient's local post office, and on to the final destination (home or office). Each transition offers each piece of contaminated mail a chance to shed some of its anthrax spores. The model distinguishes four rough categories of contamination: the original letters (of which there were six in the 2001 attack), which are assumed to have in excess of ten billion spores each, and three classes of cross-contaminated letters, with 10^3–10^4, 10^2–10^3, and 10^1–10^2 spores, respectively. (Technically, the model violates "conservation of spores," but in simulations the number of shed spores is a tiny fraction of the total supply.) The spread of anthrax is expressed for each transition as a 4×4 matrix. The entries of the matrices are estimated according to the "density" of mail handled, which is highest at regional centers.

The model incorporates estimates of the number of people exposed to contaminated mail at each stage, divided into age brackets, and the odds that exposure will lead to illness as a function of age and exposure level. Handling a piece of mail with ten, or even a thousand, anthrax spores is unlikely to make you sick, but with thousands of such letters, someone is bound to be unlucky. (The older you are, the lower the infectious dose.) Webb and Blaser's simulation of the 2001 attack indicates that the initial six anthrax letters led to 36, 432, and 5100 cross-contaminated letters in the other three categories.

The seeming precision of these numbers is, of course, an artifact of the model. The important conclusion is that cross-contamination appears to present a significant risk, even when the original letters are carefully sealed. The six letters of 2001 were taped shut, presumably in an attempt to protect the perpetrator. Next time, the attackers could deliberately amplify the effect of leakage. A "large-scale" attack with, say, 100 loosely sealed letters sent to and from nonexistent people at phony addresses (so that each letter would pass through the system twice) could cross-contaminate tens of millions of other letters and, if allowed to proceed unchecked, lead to thousands of infections.

Infectious disease mortality in the U.S., from 1900 to 1996, according to the National Center for Infectious Diseases at the Centers for Disease Control.

Flu-id Dynamics

Martin Meltzer, a health economist at the CDC, stresses the need for modelers to keep things simple. "I believe very firmly in keeping models simple enough so that you can readily identify, of all the variables going into the model, which ones are really driving the model," he says. "The more complex models you have, the less likely you're going to readily identify what's driving the whole shebang."

Meltzer's specialty is influenza, which routinely kills tens of thousands of people each year in the U.S. In particular, pandemic influenza, which occurs when the virus mutates radically, is "nature's own bioterrorist event," he says. There were three influenza pandemics in the 20th century: in 1918, 1957, and 1968. The 1918 pandemic was the worst. It killed more than 20 million people worldwide, including half a million in the U.S. (a quarter of the population here got sick). "In many ways, pandemic influenza is *the* challenge," Meltzer says.

Meltzer has developed a software package, FluAid, to help public health officials, who operate primarily at state and local levels, analyze their needs for such basics as staff and supplies to handle the next pandemic.

FluAid (described in detail at http://www2.cdc.gov/od/fluaid/) is based on a model that Meltzer and colleagues Nancy Cox and Keiji Fukuda developed to study the economic impact of an influenza pandemic. Their model distinguishes three age categories (0–19, 20–64, and 65 and older) and two risk categories (high and low) for each. Using Monte Carlo methods to deal with the uncertainties in key variables, such as the rates of illness and death for the various age and risk groups, the model produces a range of estimates for the death toll, hospitalizations, doctor visits, and workdays missed, depending on the assumed "gross attack" rate (the overall percentage of the population who get sick enough to stay home—i.e., not counting those who soldier on, spreading the flu bug to their co-workers).

The CDC researchers' model indicates that, at a gross attack rate of 15%, approximately 20 million people will feel sick enough to stay home but not sick enough to see a doctor, another 18 million will feel sick enough to see a doctor, and about 314,000 will wind up in the hospital; about 89,000 will die. If the gross attack rate is 35% (the highest rate the modelers studied), these numbers increase to 47 million, 42 million, 734,000, and 207,000, respectively. When the economic impact is measured (in particular, with standard estimates of about $1 million for the present value of people in the 0–64 age groups), the cost of a pandemic ranges from $71.3 billion for the 15% gross attack rate to $166.5 billion at 35%.

However, the modelers write, "The intent is *not* to provide 'the' estimate of impact, but rather to examine the effect of altering a number of variables." Is there a net gain, for example, either in deaths prevented or dollars saved, if vaccination resources shift from one age group to another? How sensitive are the gains to the uncertainties in the parameters? One of the key conclusions of the model is that, whatever the gross attack rate, death is far and away the costliest aspect of a flu pandemic.

Can flu preparations help in the effort against bioterrorism? Meltzer thinks they can. A proper biodefense strategy must prepare for a variety of attacks, he says. Flu models and the implementation of public health measures based on them could provide a testbed, he points out. "That provides a challenge right there: How can you use the current situation as training and practice? Perhaps the operations research people can sink their teeth into that."

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.