



Modeling the bug, the host, the world



The Bug: This negative-stained transmission electron micrograph (TEM) depicts the ultrastructural details of an influenza virus particle. Credit: Cynthia Goldsmith, CDC Public Health Image library. The Host: Credit: Leonardo DaVinci. The World: Credit: JupiterImages.

Computational Biology

CATCHES THE

FLU


BY KATHARINE MILLER



In 1918, the so-called Spanish flu killed more than 20 million people worldwide. Almost ninety years later, we're faced with the possibility of a flu pandemic that could spread even faster in this globally-connected world.

In preparation, researchers are racing to understand what makes one flu bug more infectious or more deadly than another; how best to prevent or treat influenza; and how to control a worldwide outbreak of a deadly strain. Some of these researchers are turning to an approach that was not available in 1918: computational modeling. From the molecular scale to the cellular, organismal and epidemiological, computational biologists are teaming with experimentalists to tackle tough questions about influenza and other infectious diseases.

At all scales, complexity rules the day—making influenza an appealing target for computational research. The intricacy of viruses themselves, the many interactive components of the human immune system, and the complicated networks of human interactions that lead to disease spread, all call out for integrated models that require powerful computation.

“Complexity really does matter,” says **Ira Longini, Jr, PhD**, a professor of biostatistics at the University of Washington School of Public Health with reference to his model of pandemic flu spread in the United States, “And we have the computational ability to handle it now.”





The Bug:

Modeling Shape-Shifting Viruses

The flu virus is an evolutionary marvel. Teams of experts design an appropriate flu vaccine annually just to keep up with the microbe's ability to evade the human immune system. Multiple strains circulate, and no one can predict when a new strain will emerge by mutation or recombination with another strain so that it can jump from another species to humans.

Computational biologists approach


this ever-changing bug from several angles: some simulate entire virus particles to detect their vulnerabilities; others model viral evolution to predict future strains; still others use bioinformatics approaches to design better vaccines.

MODELING THE VIRUS PARTICLE

Klaus Schulten, PhD, and colleagues at the University of Illinois at Urbana-Champaign recently simulated an entire virus particle—the satellite tobacco mosaic virus (STMV), one

of the smallest known viruses (see the News Bytes section of this issue). Allowing all of the virus's one million atoms to move for 10 nanoseconds showed surprising features of the tiny particle—and hinted at possible interventions to prevent infection.

Whole-virus simulations for flu—1000 times bigger than STMV—are still a ways off. But researchers could simulate pieces of the viral capsid—the exterior casing that holds a virus's genetic material—or they could model some parts of



The HA protein from an avian flu virus (A/Vietnam/1203/2004) with one of the molecule's three receptor binding sites highlighted as a surface representation (purple) and a close-up of the receptor binding domain with key residues labeled. In her computational models, Robin Bush has found that mutations around the HA binding pocket produce flu strains with greater long-term fitness. Courtesy: James Stevens, PhD, assistant professor in the Department of Molecular Biology at The Scripps Research Institute.

the virus in atomic-level detail while leaving other parts imprecise.

PREDICTING FLU STRAIN FITNESS

Robin Bush, PhD, associate professor of ecology and evolutionary biology at the University of California, Irvine, is modeling how specific flu virus surface proteins evolve. For flu, evolutionary fitness is largely determined by the virus's ability to evade the host's immune system.

In a 1999 paper in *Science*, Bush proposed a way to predict which of the then-current lineages of influenza A was evolutionarily most fit—that is, likely to have the most descendants.

"In H3N2 [a common strain of influenza A], we have a long skinny family tree with many lineages that quickly go extinct," she says. "Why is this?" To answer that question, Bush focused on the gene for haemagglutinin (HA), a flu virus surface protein that provokes a strong immune system response. She found that the fit strains exhibited changes in amino acids in the HA binding pocket—the place where antibodies of the immune system latch onto the flu virus.

"It doesn't take much in the way of amino acid changes to keep an antibody from binding again," she says. "Antibodies are very specific. So it's not surprising that changes around the binding pocket affect the fitness of the virus."

Bush then attempted to computationally model which mutations around the HA binding pocket would lead to long-term fitness. In 9 of 11 simulations, she found that mutations in any of 18 specific amino acids predicted that a strain's descendants would continue to infect

to pick one strain over another, she says, "you *might* pick the one that had the predicted binding pocket changes."

FILTERING THE VIRAL GENOME TO DESIGN VACCINES

Anne De Groot, MD, associate professor of medicine at Brown University, is tackling vaccine development head on. She's using bioinformatics to rationally design vaccines.

Annual flu vaccines are produced by

EpiVax fishes for antigens using epitopes as bait.

"It's a way of filtering genome information to find what's immunologically relevant," says Anne De Groot.

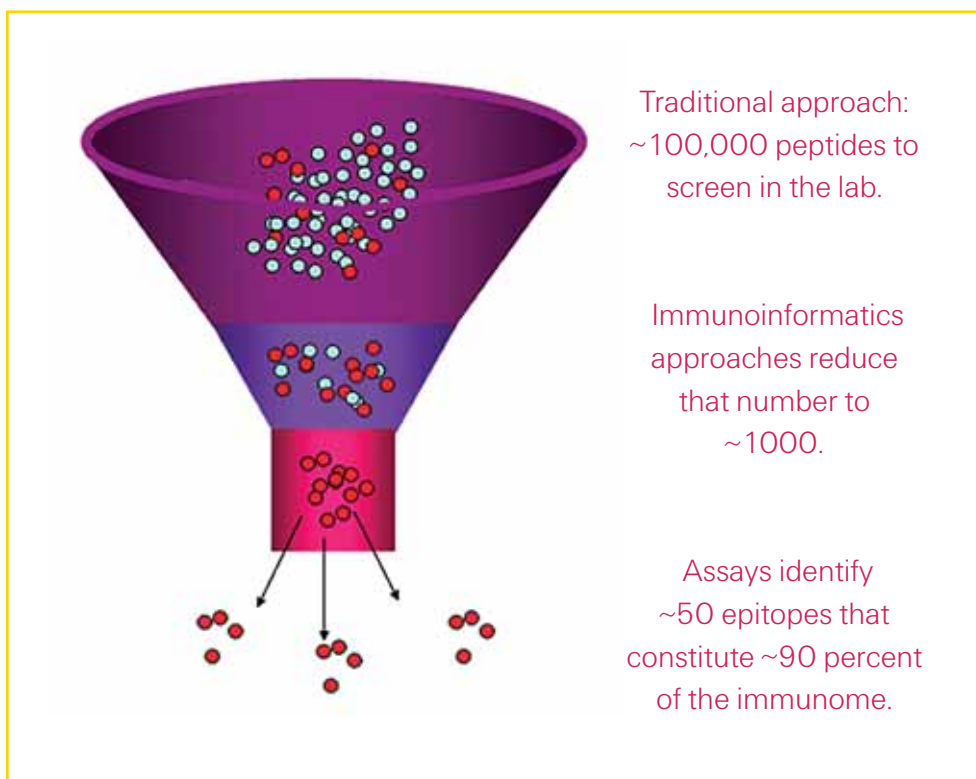
humans in ensuing years. But, Bush says, she is unable to predict if or when those expected descendants would appear.


Bush also cautions that her work is not likely to contribute greatly to annual flu vaccine design. Such vaccines contain three different flu viruses, and decisions about which lineage of each to include are made by experts based on many factors. If there was no other way

growing viruses in eggs, killing them, and then combining the dead viruses with other ingredients known as adjuvant. The process is slow, so vaccines must be designed several months before the flu season begins, with strains from the prior flu season. Vaccines containing the entire contents of dead viruses—including tens of thousands of proteins with unknown side effects—can also be

risky. "You have to be very careful about what you put in a vaccine," says De Groot, pointing to the Lyme disease vaccine that appears to have caused arthritis in some patients. "We've been lucky with some other whole-virus vaccines such as polio and cholera that have not produced deleterious effects, but all of the proteins produced by a virus have potential cross-reactivity. When you create an immune response to them, you could be

The traditional approach to epitope-mapping a typical pathogen genome could involve synthesizing 100,000 overlapping peptides. An immunoinformatics approach remarkably reduces that figure to ~1000, accelerating the discovery of ~50 epitopes that comprise over 90 percent of the immunome in an actual infection. Courtesy of Anne S. De Groot.





creating auto-immunity or pre-setting an immune response that you don't want."

A different approach is to put the vaccine together one piece at a time, so you know exactly what's going on, De Groot says. In addition to being safer, this approach should allow development of a vaccine in response to the current flu strain (rather than last year's) because individual pieces (peptides) can be rapidly manufactured.

EpiVax, a Rhode Island biotech company founded by De Groot in 1998, uses computational tools to design peptide-based vaccines. They use a process called fishing for antigens using epitopes as bait. "It's a way of filtering genome information to find what's immunologically relevant," says De Groot. In the 1990s, researchers developed algorithms that can pick out gene motifs that are likely to stimulate the immune system. Using their own version of such algorithms, known as EpiMatrix, EpiVax filters a particular pathogen's genome to pick out snippets likely to produce immuno-stimulatory peptides known as epitopes. These peptides can then be synthesized in a lab and mixed with blood from people who have been previously exposed to the particular pathogen. The epitopes that successfully "fish out" responses in the blood are presumably part of an antigen—one of the viral or bacterial proteins to which the person's immune system responded during the earlier infection. Such antigens and/or their epitopes are potential ingredients in a vaccine, since they produce valuable immune responses. This approach has led to potential vaccines for HIV and meningitis that are now in clinical trials.

A bioinformatics approach might also contribute to development of a universal flu vaccine, De Groot says. Algorithms can screen the genomes of all the various flu strains to look for genomic sections that are pretty short and don't change. "They're kind of like the flu thumb or index finger: they are critically important to the function of the virus," De Groot says. Running these regions through another algorithm will reveal whether they stimulate the immune system. If they do, then a flu vaccine containing these proteins might induce immunity to a group of flu strains rather than just one.



The Host: Modeling the Immune System

Using computation to understand the flu virus and its proteins only covers half the story. The T-cell mapping interface used by De Groot hints at the other half: the host immune system.

When a virus or bacterium invades the human body, it stimulates a cascade of immune system events to fend off the intruder. Over the last hundred years, experimentalists have cleverly studied these events in contexts where only one component changes at a time. It's work

tional immunology comes from HIV work published in 1995 by **Alan Perelson, PhD**, and **David Ho, MD**. It led directly to the realization that HIV could be treated with cocktails of drugs—an approach that has greatly reduced the number of deaths due to AIDS.

This research demonstrated that it's not too soon to take computational immunology seriously, Kepler says. Moreover, he adds, "the rate of accumulation of new information is so fast, that if

"Computation is a way to take all these objects [the pieces of the immune system] and put them back together into a form where the goal is not to minimize variation but to keep track of it," says Thomas Kepler.

that has generated huge amounts of data about more than 20 different types of immune cells and a few thousand participating molecules. But what's missing, say computational immunologists, is an integrated view of the puzzle.

"Computation is a way to take all these objects and put them back together into a form where the goal is not to minimize variation but to keep track of it," says **Thomas Kepler, PhD**, professor of biostatistics and bioinformatics at Duke University.

The current poster-child of computa-

we don't start now, we'll never catch up."

It's a view shared by leaders at The National Institute of Allergy and Infectious Diseases (NIAID) who, in 2004 and 2005, funded four computational immunology projects. Three of these are using flu as a model pathogen.

THE IMMUNE SYSTEM AS A BLACK BOX

Under an NIAID grant to **Penelope Morel, MD**, associate professor of immunology at the University of Pittsburgh, researchers are modeling

how respiratory infections (influenza, tuberculosis and tularemia) affect the local immune response in the lungs. The group will be gathering data about how macrophages in the lung respond to each virus by measuring such things as secretions (cytokines) and cell surface-markers as they change through time. But the goal is to take the experimental measurements and plug them into computational models. “If your model doesn’t match the data, then you know something’s missing,” says Morel. “That exercise is a highly valuable one.”

The project’s flu modeler, **Shlomo Ta’asan, PhD**, professor of mathematical sciences at Carnegie Mellon University, is taking a highly mathematical approach: He will look at the immune system as a black box, without

“The model is easy to write out,” says Hulin Wu, “But there’s no validation without data.”

making assumptions about the biology.

“We don’t put anything into the model except the data that come out of the experiments,” he says. “Our algorithm will spit out something that might be intuitive for biologists, and it might not.” He hopes to find out if math can cut through biological intuition to gain

some new truth. After creating a model that seems to reproduce the experimental results for macrophage responses, Ta’asan says, “Then we want to see how to manipulate it with various drugs.”

The biggest challenges to Ta’asan’s model are practical ones. One mouse doesn’t have enough blood to cover all the necessary tests and must be sacrificed to get certain measurements. In addition, microarray data are highly variable and there is fuzziness in the measurements. Ta’asan says some people simply ignore that variability, but he thinks it says a lot about the system and should be accounted for mathematically. “We’re thinking about using some fuzzy logic ideas or probabilistic approaches,” he says. “We don’t want to pretend it’s not a problem.”



This 1976 photograph shows an elderly female receiving a vaccination by a public health clinician during the nationwide swine flu vaccination campaign, which began October 1, 1976. Courtesy: Centers for Disease Control and Prevention.

MODELING THE IMMUNE SYSTEM USING EXPERTISE

Hulin Wu, PhD, professor of biostatistics and computational biology at the University of Rochester shares these concerns. He, like Ta'asan, received an NIAID grant and is modeling the immune system response to flu. But Wu is taking a more traditional approach: He develops his models based on immunologists' and virologists' current theories about flu infection. And he needs lots of data on the kinetics of the virus and the cells with which it interacts. For example, he needs to know how fast the flu virus proliferates and dies; the infection rate for various cell types; and the rate of production of T-cells, antibodies, CD4 and CD8 cells, and lymphocytes. On top of that, he needs this data from several different locations (lung and lymph nodes, for example) at various time points so that he can model the host

“It’s difficult to understand parallel events without the benefit of computational approaches,” says **Stuart Sealfon**.

reaction as the virus and immune cells migrate between compartments.

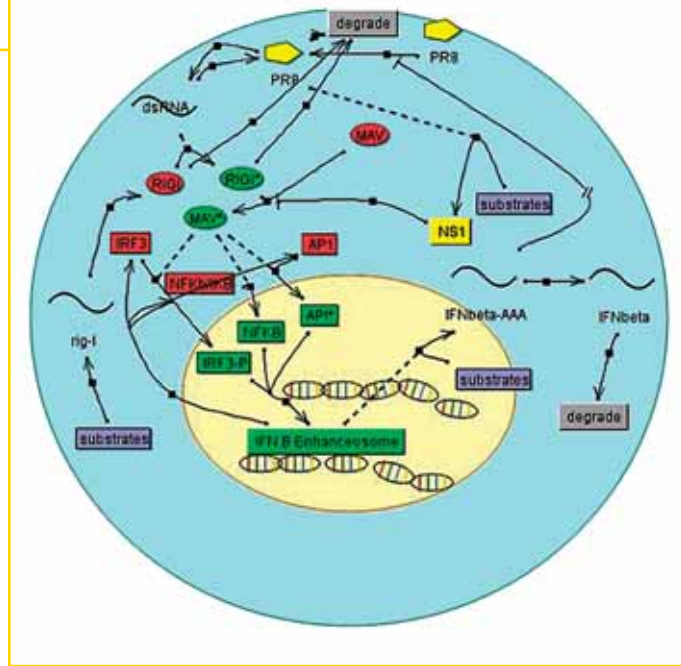
But he’s finding that such data just doesn’t exist for flu. Coming from HIV modeling, this can be frustrating. “HIV is a long-term infection. You can measure the immune response over many years,” he says. “Flu lasts only one week, and then everything’s gone.” Getting enough measurements in a short time span is challenging but essential. “The model is easy to write out—to describe the interactions between the virus and the immune system in the lung, the spleen and the lymph nodes. But there’s no validation without data.”

The data-gathering problem would be even worse in the event of a bioterrorist event, he says. “How can we collect enough information quickly to deal with a new engineered virus?” That’s when an immune system model would prove valuable. If there’s a model in place for an existing flu virus, it can be quickly adjusted to a new one, he says.

MODELING MOLECULAR LEVEL IMMUNE RESPONSES

Another NIAID group led by **Stuart Sealfon, MD**, professor of neurology at Mount Sinai School of Medicine in New York City, is using computation to get a handle on the immune system’s response to flu at the molecular level. They are modeling the ways that flu viruses evade or undercut the immune system’s efforts, specifically focused on the dendritic cell—the transitional cell between the innate and adaptive immune systems.

The team starts with experimental work: They infect dendritic cells with non-pathogenic viruses containing specific components of the flu virus such as NS1 (a protein that shuts down some parts of






This is the graph of a preliminary 18-equation model of flu virus antagonist effects on interferon production in dendritic cells. Both the graph and model were developed by Mount Sinai researchers using BioPathwise, a signaling simulation program developed by BioAnalytics Group LLC. Courtesy of Stuart Sealfon.

the normal signaling in such cells). This generates large amounts of data on gene and protein changes. The computer model then tracks all of these changes at once. “It’s difficult to understand parallel events without the benefit of computational approaches,” Sealfon says.

One of the modeling challenges, Sealfon says, is dealing with events that occur on different time scales. Signaling events take place over minutes, gene induction occurs over hours or a few days, and secretion and stimulation occur throughout the infection period. These multi-scale modeling problems still need to be addressed, he says. But if the challenges can be overcome, “ultimately, this work can help us to develop strategies to circumvent the virus’s actions.” And in the event of a new strain, the model can help identify the evasive tactics used by the new flu bug, which might lead to an appropriate therapy or vaccine.

Computational immunology still has a long way to go before it will fulfill its promise, Kepler concedes. But the field is really opening up, as technology provides more and more ways to measure the many complex interactions of the immune system. “There has already been a lot of good work in computational immunology,” he says, “but it will have a very different flavor in the next few years.”





The World: Modeling Flu Spread

The field of computational epidemiology is a much more mature field than computational immunology, Kepler says. Because epidemiologists have always dealt with disease spread across large populations, it's not as big a leap to computation on a national and global

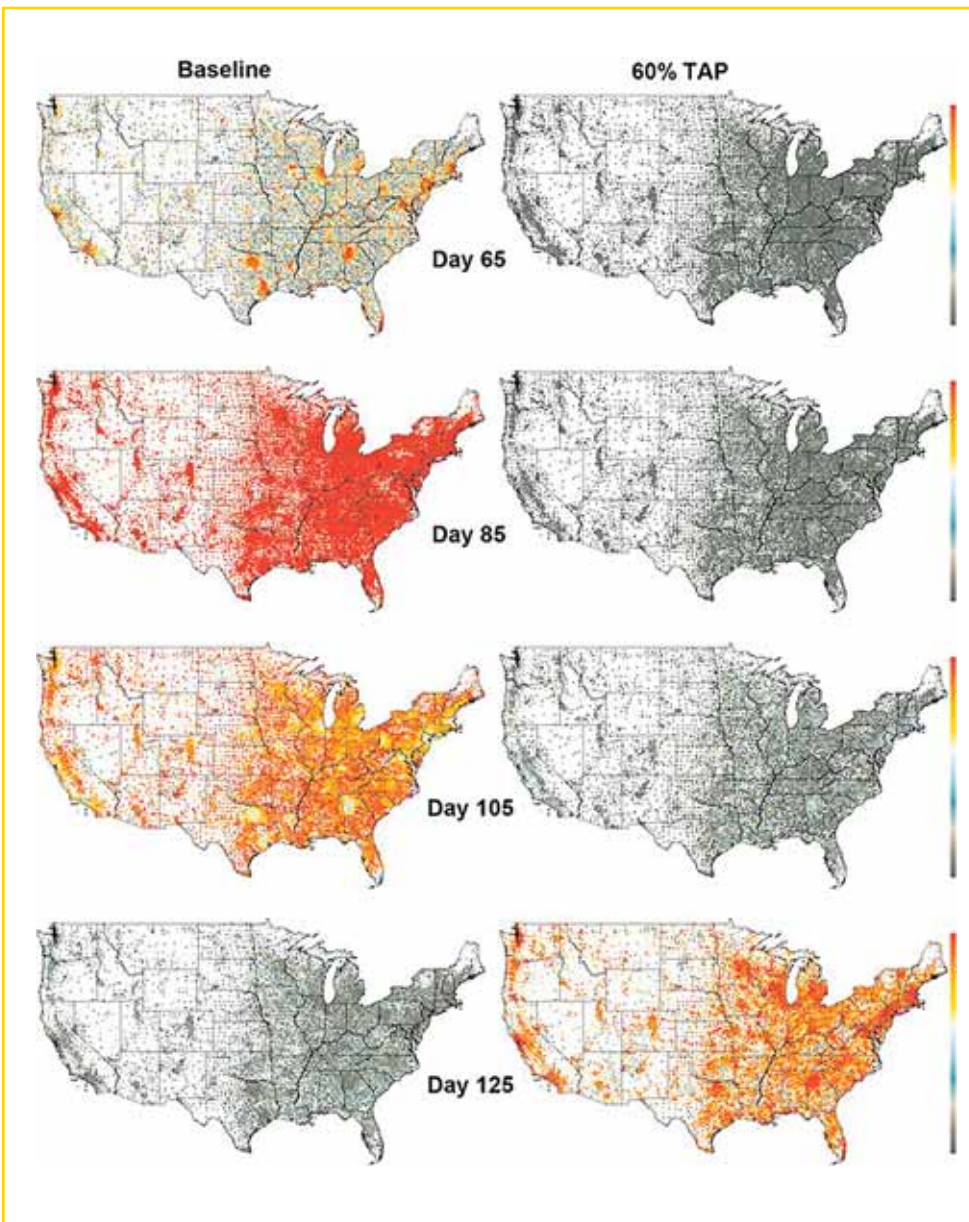
level. And the main computational approaches to epidemiological problems—agent-based modeling and graph theoretical methods—are well-established.

What is new, however, is the current United States effort to bring infectious disease modeling under one umbrella. In

2004, the National Institute of General Medical Sciences (NIGMS) within the NIH created the Modeling of Infectious Disease Agent Study (MIDAS), a program that funds several epidemiologic modeling efforts, gives them access to supercomputers, and also coordinates them in hopes of producing results that will be useful to policymakers.

MIDAS literally gets everybody in the room—programmers, data collectors, database designers, biologists, epidemiologists and statisticians—to try to iron out all the potential areas of disagreement. In particular, they try to reach consensus about what parameters should be part of the model. Telling policymakers that one program

On day zero in these simulated pandemic outbreaks (R_0 of 1.9), infected individuals arrived at 14 major international airports in the continental United States. The number of ill people at a given point in time is indicated by the color scale from green (0.3 percent) to red (3 percent) of the population. More than 40 percent of the entire U.S. population end up getting ill with no interventions (left). The use of antivirals slows the spread (right) until the stockpile of 20 million courses runs out, at which point there's a delayed nationwide pandemic. Courtesy: PNAS.





gets one result and another program gets a different result simply won't do, says **Irene Eckstrand, PhD**, scientific director for the MIDAS program. "So we try to work all those things out in-house."

The end goal is for MIDAS to be able to tell policy makers: Based on our models, a specific intervention in a specific type of epidemic will likely have a specific effect. But, Eckstrand cautions, the models are all stochastic—they don't give the exact same answer back twice. Uncertainties are built into the models because many parameters are probabilistic. For example, the likelihood that a person will stay home on a given date rather than spread the disease to one or more people can be assigned a specific probability so that the outcome will vary each time the model runs. So each computer model must be run multiple times on a given set of parameters in order to produce a distribution of results that express the range of possible outcomes as well as the most likely outcomes.

Although the MIDAS approach could be applied to any infectious disease, the researchers decided early on—before the current concern over avian flu—that it would be interesting to model pandemic influenza. "The timing was pretty remarkable," says Eckstrand.

Because of this fortuity, MIDAS models published in *Science* and *Nature* in August 2005 and in *Proceedings of the National Academy of Sciences (PNAS)* and *Nature* in April 2006 were front page news.

MIDAS grantee **Ira Longini** co-authored two of these high-profile papers. His August 2005 paper in *Science* looked at ways to stop an outbreak of flu in an imaginary population of 500,000 people in Southeast Asia. He and his colleagues found that an outbreak could be contained if a sufficient stockpile of antiviral drugs could be delivered rapidly enough—within three weeks of the first human-to-human transmissions. In practice such an

approach would be difficult to implement in Southeast Asia but the model will help policymakers plan for an effective response.

MODELING FLU ACROSS THE UNITED STATES

Longini's April, 2006 model in *PNAS* focused closer to home: What interventions would help contain a flu pandemic in the United States? Instead of an imaginary population, this model was built on census tract data for 281 million people and relied on extensive knowledge about peoples' travel and activity patterns. "We're all pretty predictable, really," he says. "We all get up, go to work, go shopping, and get together

"These models aren't meant to be predictive tools," Longini says. "They are meant to evaluate strategies for intervention."

with our neighbors." So Longini's model breaks down social contacts into 12-hour time periods (day and night) in seven different contexts ("mixing groups"). In some contexts, close contact occurs (home, work, schools); in others, it's more occasional (shopping malls).

The key variable for Longini's model is a number called R_0 , which represents how transmissible a strain will be. Specifically: R_0 is the number of people, on average, that a typical infectious person infects during the infectious period in a fully susceptible population. If that number is bigger than one, then the disease will spread. Less than one and it will disappear.

No one really knows what the R_0 for a new pandemic flu strain would be. It's thought that newly emerging strains that haven't had a chance to adapt to humans might have a low R_0 and there-

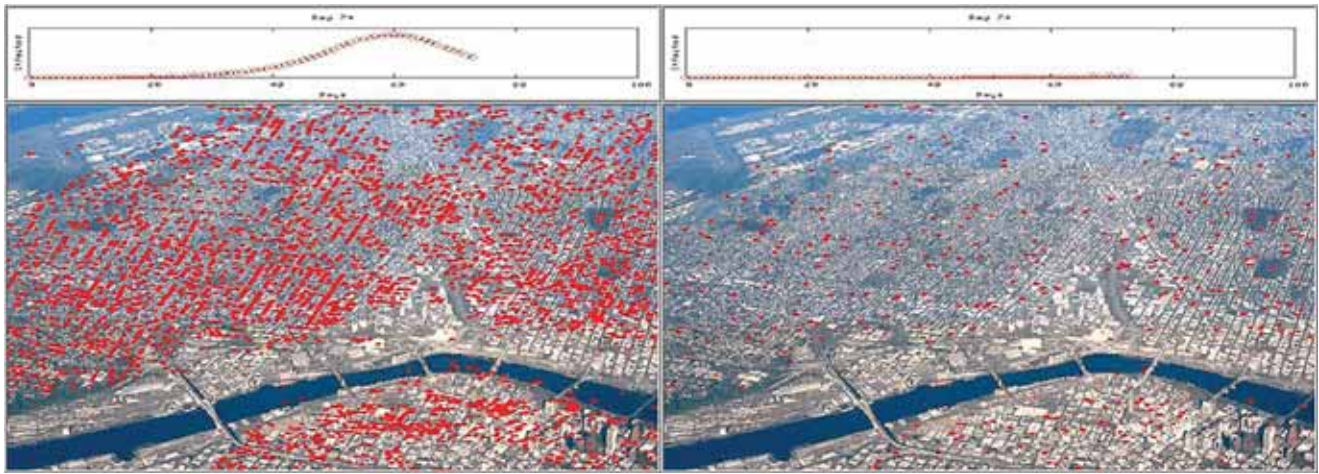
fore may die out. But no one has observed an emerging infectious disease before it becomes well adapted. "We kind of missed HIV and SARS," Longini says. But now, with better surveillance, virology and field epidemiology, "Flu might be the first emerging disease where we really have an opportunity to watch what happens."

In Longini's computer model of the United States, he's assuming a well-adapted virus, so he starts with a pretty high R_0 of 1.6 to 3.0 (the R_0 for Smallpox is 5; for the 1918 flu, about 2). But R_0 is only the starting point for the model. As different intervention strategies are tried, the R value changes. "These models aren't meant to be predictive tools," he says.

"They are meant to evaluate strategies for intervention."

For a flu pandemic with an R_0 of 1.6, Longini and his colleagues found that any of several individual strategies such as antiviral drugs, child-first vaccinations or school closures could be fairly effective in reducing the incidence of flu below ten percent (the rate for a typical annual flu season). If the R_0 is higher than 1.9, however, only vigorous application of multiple strategies would reduce the outbreak's impact.

Longini has been working directly with the government on intervention scenarios. For example, he can compare the impact of stockpiling 10 million versus 100 million courses of Tamiflu. And he can say closing schools is more effective than other social distancing measures while travel restrictions appear to have little impact (findings confirmed by the other MIDAS model for the United States published in *Nature*). But he's quick to point out that the model cannot predict what will actually happen. "We can say one strategy might be better than another or one might be totally ineffective and another has a good chance of being effective. So we can make those sorts of statements, and that's about as far as we can go."



An aerial view of Portland, Oregon, in a simulated epidemic performed by Eubank and colleagues. In each image, red dots represent a location with at least one infected person present at a set time on a given day. On the left is the baseline spread after 60 days. At right is what happens if 75% of households make the decision to stay home during those 60 days. These simulations are not meant to represent what would happen at a particular address, but to indicate the general consequences of such extreme behavioral modification. Courtesy: Stephen Eubank.

GETTING DOWN IN THE WEEDS

Despite the unifying influence of MIDAS, its grantees still have their own particular approaches to modeling epidemics, says Eckstrand. “There’s an interesting discussion about how much detail you need to know in order to build higher level estimates,” she says. **Stephen Eubank, PhD**, is a MIDAS grantee who works with “down in the weeds” information, Eckstrand says.

Eubank, project director of the Network Dynamics and Simulation Science Laboratory at the Virginia Bioinformatics Institute, believes models need detail in order to best address policymakers’ needs. To evaluate the relative effectiveness of strategies such as telecommuting, limiting meeting sizes or setting quotas on the number of people in a grocery store, a model must contain sufficient details about what individuals are actually doing and where. Eubank’s model consists of individual agents that each represent a single person assigned a set of activities at reasonable locations given where they live. “So we don’t have a knob in our model that says ‘reduce contact rates by thirty percent,’” he says. “Instead we have knobs that say, keep some people home from work; or don’t let more than ten people in this room.”

Right now, Eubank’s models can only be applied to one city at a time. “It’s hard to support both the amount of detail that we’re talking about in our model and the scale of the whole country. It becomes a question of computer

resources.” For a city of a million or so, each engaged in five to ten activities, a simulation covering 60 days can take an hour or two on 30 CPUs. But expanding such detailed models nationwide would require very large clusters of computers and large quantities of data. Eubank’s hope is to develop grid-based platforms. “It’s unlikely that any one person or organization would want to model this much detail for the whole United States,” he says. “But at the local level, there are good arguments for why an urban area should have such a model of itself.” It would be useful not only in the event of an epidemic, but for other kinds of urban planning.

If cities participate in Eubank’s plan, they could then tie their models together in a grid to create a nationwide, detailed model. “So we’d have this loose federation of urban or regional models interacting across the grid, each maintained by someone with a vested interest in having a good model of their area.”

OPTIMIZING INTERVENTIONS

Catherine Dibble, PhD, assistant professor in the department of geography at the University of Maryland, College Park, also offers a different perspective within MIDAS. As a collaborator on the MIDAS grant headed up by **Donald Burke, PhD**, at Johns Hopkins University, she has developed tools for


doing two things many others haven’t done: risk analysis and optimization. “Most pandemic modelers decide the interventions and settings by hand and run them through the simulations,” she says. “We do that too, but we can also optimize interventions and evaluate their risks.”

So, while Longini and other MIDAS


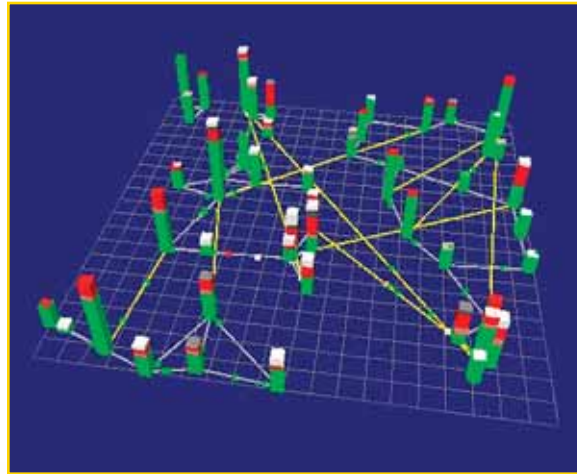
Telling policymakers that one program gets one result and another program gets a different result simply won’t do, says Irene Eckstrand, “So we try to work all those things out in-house.”

modelers (such as **Neil Ferguson, PhD** and **Mark Lipsitch, PhD**; see www.epi-models.org) recommend which local interventions and combinations of interventions could be most effective, Dibble has the capacity to evaluate the optimal geographic deployment of those recommended interventions and associated scarce resources such as Tamiflu and vaccine supplies.

In addition, Dibble’s risk analysis tools can evaluate the optimal strategies



Catherine Dibble's approach to flu modeling is similar to her earlier work with SARS (published in 2003). In this simulated landscape of SARS spread, each tower represents a population center. Colors represent the health status of individuals at each center:
 green = healthy
 pink = infectious
 red = symptomatic
 gray = deceased
 white = recovered (now immune).



“If these models can be useful at all,” Catherine Dibble says, “people need to be comfortable with them and understand how a particular intervention can help.”

to see how well they deal with events that don’t go as planned. As Dibble explains, “some interventions might give a good outcome under some conditions, but, compared to other possible interventions, might be more sensitive to chance events that could work against them.”

But optimization and risk analysis would require huge amounts of computer resources if applied to the fully detailed national models, Dibble says. “Effective optimization requires a model that represents key aspects of geographic structure and travel behavior, yet is simple enough to run hundreds of thousands of times to fully explore alternative geographic deployments and to explore uncertainties, sensitivities and risks.”

Dibble’s model is designed to evaluate the effect of travel restrictions between transportation hubs in the

event of pandemic flu. She created a network with healthy individual agents (green) distributed at each transportation hub all across the continental United States. Each population center can be visualized as a tower with its height determined by its relative population. “Then we drop one or more infected individuals into the landscape,” says Dibble. “They are pink.”

As time goes by, different people make different travel decisions (modeled using actual airline routes and travel data) and the infected agents start “sneezing” on people (infecting them) at a rate consistent with a particular R_0 and whichever interventions may be imposed. Sometimes the epidemic fizzles out—the equivalent of the infected person going home and not giving the disease to anyone. When it doesn’t fizzle out, pink

(infectious), red (sick), gray (dead), and white (recovered) people appear on the landscape, with travel decisions leading to diffusion among cities. “We focus on evaluating the relative pandemic risks across cities: Which cities in the United States are likely to be hit soonest or more often,” she explains.

In the event of a pandemic, her model can suggest how to allocate the available (and limited) resources effectively, Dibble says. Spreading interventions uniformly over the population might seem fair, but it might not control the pandemic as effectively as targeting the resources to particular cities.

Convincing policymakers to focus resources geographically could be a big challenge, she says.

“If these models can be useful at all, people need to be comfortable with them and understand how a particular intervention can help.” That kind of public awareness, she says, will be key. According to her, “Communication may turn out to be more important than any particular model, vaccine or resource.”

BRINGING BUG, HOST AND WORLD TOGETHER

As with many modeling endeavors, the question arises: What if the models could be integrated across the scales? Will we eventually model an evolving flu virus interacting with the host immune system in such a way as to predict, with reasonable reliability, its effect on a population? If so, that day is not near. But even now, efforts by MIDAS researchers might help stem the spread of a flu pandemic, potentially saving millions of lives. Even if the models don’t help, and a pandemic rages uncontrollably, the work will help prepare us for the next time, and the one after that. “Pandemic flu is a big threat,” says Longini, “but it’s also a really important scientific opportunity.” □