

WASHINGTON ROUNDTABLE
ON SCIENCE & PUBLIC POLICY

**SECURITY IN SOCIETY:
Protecting an Increasingly
Connected World**

By Harvey Rubin

GEORGE C.
Marshall
I N S T I T U T E

Washington, D.C.

The George C. Marshall Institute

The George C. Marshall Institute, a nonprofit research group founded in 1984, is dedicated to fostering and preserving the integrity of science in the policy process. The Institute conducts technical assessments of scientific developments with a major impact on public policy and communicates the results of its analyses to the press, Congress and the public in clear, readily understandable language. The Institute differs from other think tanks in its exclusive focus on areas of scientific importance, as well as a Board whose composition reflects a high level of scientific credibility and technical expertise. Its emphasis is public policy and national security issues primarily involving the physical sciences, in particular the areas of missile defense and global climate change.

The Washington Roundtable on Science and Public Policy

The Washington Roundtable on Science and Public policy is a program of the George C. Marshall Institute. The Roundtable examines scientific questions that have a significant impact on public policy and seeks to enhance the quality of the debate on the growing number of policy decisions that look to science for their resolution.

The opinions expressed during Roundtable discussions do not necessarily represent those of the Marshall Institute or its Board of Directors. Additional copies of this transcript may be ordered by sending \$7.00 postage paid to:

The George Marshall Institute
1625 K Street, NW Suite 1050
Washington, D.C. 20006
Phone: 202/296-9655
Fax: 202/296-9714
E-mail: [info @marshall.org](mailto:info@marshall.org)
Website: www.marshall.org

Security in Society:
Protecting an Increasingly Connected World

By Harvey Rubin

The George Marshall Institute
Washington, D.C.

Harvey Rubin is a member of the Marshall Institute's Bioterrorism Steering Group and is the Director of the Institute for Strategic Threat Analysis and Response (ISTAR) at the University of Pennsylvania. He is also a professor of medicine, of microbiology and of computer science at that University.

SECURITY IN SOCIETY

*Protecting an Increasingly Connected World**

Harvey Rubin
April 29, 2003

Summary

From airline security to the internet to vaccine distribution, our society is growing increasingly interconnected. This connectivity has brought great benefits, but has also increased our vulnerability. Society now has many points of access vulnerable to attack and greater interdependence has made it difficult to separate or secure any particular area within the system. Existing approaches to the problem of system vulnerability do not adequately assess how risk is shared among participants or determine the best strategies to protect them. In the case of vaccination, for example, each person's decision to be vaccinated or not has associated implications and costs which can affect the wider society as well as than the individual. Dr. Rubin argues that key members within the social network – medical professionals, government and community leaders – must take the lead in developing efficient strategies to deal with threats to the complex infrastructure systems which support the industrialized democracies.

Dr. Rubin's Remarks

Thank you very much. I chose a topic that we are working on at the University of Pennsylvania called Interdependent Security of Complex Networks. We are going to talk about a little bit of policy, some mathematics, some economics, some biology and hopefully we will be able to tie it all together.

What is a network and what is the interdependent security of complex networks? Let's define, at least for the purposes of today, a network as the relationship between individuals or groups that is generally mutually beneficial, stable, robust and secure. There are times when we want to keep the network intact, but there are times when we want to destroy networks as well, so we have to think about both sides, about maintaining the network as well as breaking into the network. A catastrophic state of af-

* The views expressed by the author are solely those of the author and may not represent those of any institution with which the author is affiliated.

fairs could occur when interdependent security within the network fails. This can happen either by design or inadvertently, and it can be either anticipated or not anticipated. Determining whether the failure happens by deliberate attacks or random failures is a very important distinction to make, and I will tell you why, even mathematically, we distinguish between an attack on a network and a failure of a network. Those are two fundamentally different physical and mathematical aspects of the network, with different processes involved in each. To do this, we need to define our terms, provide some examples, show how these terms can actually be formalized and analyzed mathematically and then we will try to pose some questions, the answers to which constitute a very deep and rich research process.

I will make a short comment about the University of Pennsylvania's Institute for Strategic Threat Analysis and Response (ISTAR) that was started a little over a year ago. The goal of the Institute is global; we are taking on issues that are both intentional and unintentional in terms of threat. The Institute is composed of faculty from all twelve Schools at Penn. We do not do classified research at Penn, so everything we do is publishable and public. Some of the areas we are working on include: the dynamics of natural and social networks; control of biological agents such as pox viruses; health care delivery and public health systems, especially as it relates to bioterrorism; protecting the information infrastructure; modeling transportation systems; medical, nursing, veterinary and dental curriculum development to teach the response to all hazards; economic models of regions, states and nations; and legal issues and issues of criminal justice. Governor Ed Rendell just appointed Jim Eisenhower Chairman of the Pennsylvania Commission on Crime and Delinquency and one of his missions is to expand into areas of Homeland Security. I am fortunate to serve on that committee with him.

What is interdependent security? Let's take some real examples. The first example comes from my colleague at the Wharton School at Penn, Dr. Howard Kunreuther, and Dr. Geoff Heal at Columbia University. Suppose you are an airline, Airline A, you have a certain incentive to protect your baggage, so you invest an enormous amount of money in baggage security. Airlines B, C, D etc. say, "Well, gee, we should spend a lot of money to protect our planes and customers too". The whole system becomes very secure. However, there is less incentive for airlines to invest in baggage security when one airline does NOT invest in security. Why? Because once checked in, carriers receiving baggage in transit from other airlines may not re-screen the baggage and the dangerous bag will get through

whether or not there is a security system in place for newly checked in customers. Where do I put my money? Do I put a lot of money on airline security or not? So you can imagine starting to play that game, which is a multi-agent game on a network that has structure. It becomes very complicated to calculate the cost benefits and equilibrium points of the system. Our colleague Michael Kearns and his students in the Computer Science at Penn are working out some of the beautiful mathematics of these systems.

This is what we have been calling interdependent security. My decision to act depends not just on the transaction, but it depends on what the other actors decide to do as well and it depends on the structure of the network that the other actors comprise. Take another example: vaccination strategy. Forget about smallpox for a moment and think about the serious and present danger of influenza, which still kills tens of thousands of people in the United States each year. Influenza starts off in an animal community somewhere and spreads around. In the United States, children become an important repository of influenza. Who gets very ill and dies? People over 65 and people with some chronic diseases. Here are the recommendations for vaccination policy as they stand today:

Vaccination is recommended for the following groups of persons who are at increased risk for complications from influenza:

- persons aged ≥ 65 years;
- residents of nursing homes and other chronic-care facilities that house persons of any age who have chronic medical conditions;
- adults and children who have chronic disorders of the pulmonary or cardiovascular systems, including asthma;
- adults and children who have required regular medical follow-up or hospitalization during the preceding year because of chronic metabolic diseases (including diabetes mellitus), renal dysfunction, hemoglobinopathies, or immunosuppression (including immunosuppression caused by medications or by human immunodeficiency [HIV] virus);
- children and adolescents (aged 6 months--18 years) who are receiving long-term aspirin therapy and, therefore, might be at risk for developing Reye syndrome after influenza infection; and
- women who will be in the second or third trimester of pregnancy during the influenza season.
- physicians, nurses, and other personnel in both hospital and outpatient-care settings, including medical emergency re-

- sponse workers (e.g., paramedics and emergency medical technicians);
- employees of nursing homes and chronic-care facilities who have contact with patients or residents;
 - employees of assisted living and other residences for persons in groups at high risk;
 - persons who provide home care to persons in groups at high risk; and
 - household members (including children) of persons in groups at high risk.

"Because children aged 6--23 months are at substantially increased risk for influenza-related hospitalizations, influenza vaccination of all children in this age group is encouraged when feasible. However, before a full recommendation to annually vaccinate all children aged 6--23 months can be made, ACIP, the American Academy of Pediatrics, and the American Academy of Family Physicians recognize that certain key concerns must be addressed. These concerns include increasing efforts to educate parents and providers regarding the impact of influenza and the potential benefits and risks of vaccination among young children, clarification of practical strategies for annual vaccination of children, certain ones of whom will require two doses within the same season, and reimbursement for vaccination. ACIP will provide updated information as these concerns are addressed. A full recommendation could be made by 2003--2005. In the interim, ACIP continues to strongly recommend influenza vaccination of adults and children aged ≥ 6 months who have high-risk medical conditions.

The current inactivated influenza vaccine is not approved by FDA for use among children aged <6 months, the pediatric group at greatest risk for influenza-related complications. Vaccinating their household contacts and out-of-home caretakers might decrease the probability of influenza among these children. " (Prevention and Control of Influenza Recommendations of the Advisory Committee on Immunization Practices (ACIP)
<http://www.cdc.gov/mmwr/preview/mmwrhtml/rr5103a1.htm#RecVac>)

This strategy has worked relatively well. But consider the network: you go to put your child in day-care, then you go see Grandma and Grandpa in the assisted-care living facility, then you go to work. This is your network and you're the link, the link between the children who may

have influenza and are not vaccinated and the folks in the nursing home where people are going to die from influenza, and then you go back to work and, since you may not have been vaccinated, you give influenza to your colleagues and they go to their networks and it grows and grows and grows. What is the cost? The cost is getting sick and dying, if you are over 65, or possibly having a reaction to the vaccine. What is the benefit? Well, if you vaccinate the over-65s, they will survive and they won't die of influenza or related infections. We could have a perfectly secure system if the at-risk people are vaccinated along with some, but not necessarily all "links" between certain populations. And what happens in the event of an outbreak of a pandemic? In all cases, we can ask, "where is the Nash equilibrium established?" It turns out that the analysis of this question involves some of the same kind of mathematics as the analysis of the airline security issue.

There are a number of similar systems, such as baggage security, computer systems, and the vaccination story. There is a very famous story that is described in Duncan Watts' recent book *Six Degrees*, which I recommend to you, and two others as well, one by Albert-László Barabási called *Linked* and another by Steven Strogatz called *Sync*. Watts gives the example of the Toyota brake manufacturing company. You see, Toyota made all of a key component of its brakes in one company, the Aisin Brake Company, and that place burned to the ground. Because they had an incentive to keep very low inventory, they only had two or three days worth of brake parts. Aisin burned to the ground and they had literally three days worth of parts, and without any direct order from the higher-ups at Toyota, the subunits of Toyota managed to reconfigure themselves and supply brake parts and within a week, they were back up to almost full scale production of cars. According to Watts and the references he cites, how that happened is a wonderful example of networks working well. This was a failure, not an attack, of a complex network, and a really interesting example of how the system, the network, basically healed itself. Learning how to *design and build* self-healing networks is a wonderful and terribly important research effort.

It is clear that networks are susceptible to attack; we agree on that. But can we quantify that? Is there some way to say that I am more susceptible or less susceptible? And if we could do that, could we think of the best place to attack, if we might want to attack a network? And are all networks the same? We need to define a few things. What is a network? We are going to define a network as a collection of individuals. In the literature they are called actors, agents, components, nodes, vertices; I like to call

them nodes, but sociologists call them actors. These nodes are connected by links, or edges. That link could be simply *knowing* someone or it might be having *influence* over someone, or going to school together. Networks can be directed, that is, I can be the father of my son, but he is not the father of me, so the link goes in one direction, or the network might be not directed, that is, we just know each other and are basically connected that way. When you build up structures of networks, the direction of links becomes that much more important. We see this in the work I do in the lab analyzing metabolic and genetic networks where biochemical reactions flow in a particular pathway and regulatory elements act in a particular and often very specific way. Links can have value, so that I can be linked to you, even if I just barely know you, or I could be your father, in which case I have a stronger link to you. If we are companies or airlines, we can put a value on exchanges of information, so we can start building up very complicated ideas of networks.

The degree of a node is essentially the number of links it has, so my degree now has been increased by the number of people in this room, since I now know all of you. Obviously a network that is built up of many, many nodes looks like a horrible mess of spaghetti. The shortest distance, the shortest path between any two nodes is what we call a geodesic. This is where the famous “Six Degrees of Kevin Bacon” game comes in. It is actually true globally: there is approximately six degrees of separation between me and a farmer in Afghanistan. Now how can that possibly be? Well, I know you, and you, and since we are sitting in the Army & Navy Club in Washington D. C., I am sure one of you knows President Bush, and President Bush knows the Afghan President Hamid Karzai and Karzai knows somebody else in town and that person knows the farmer and there you are, six degrees. So when you think about it, it makes some sense as a back-of-the-envelope calculation. Suppose that I know a hundred people and each of those one hundred people knows another hundred people, independently, and those hundred people know another hundred people. Well, it’s a hundred times a hundred times a hundred, six times. How many people are there in the world, six billion or so? So you can see that logically it comes out to about six degrees, give or take a factor. The connectivity of the internet is about nineteen, that is, you can get from one node to any place else on the internet with about nineteen clicks.

Now there is another concept in network theory called closeness. Closeness is the distance from one node to all the other nodes, so that a node that has a high degree of closeness receives a lot of information, but also has a chance of getting a lot of contagion. So if I am highly con-

nected, that means I am going to hear about everything but I can also catch a lot of things. You can see clearly why this is important in terms of public health and people spreading sexually transmitted diseases and you can see how this serves a role in all sorts of social networks. Somebody who is really close to everybody is going to get something.

“Betweenness” is another property of networks, relating to the function of being a conduit for information throughout the network. If I am a conduit, then I have the ability to regulate that information. If I decide to slow down, and everybody has to go through me to get to somebody else, then I can control information. That is the notion of “betweenness” and that is obviously very important. If you can detect somebody who has the highest “betweenness” and the highest closeness and vaccinate them, or eliminate them if they are part of a terrorist organization, or promote them if they are in a corporation, then you’ve got the right person. You can also assign a number, or a value to these things. These are absolutely quantifiable ideas based on networks; it is basically counting. There are software programs that analyze your network: you can click on it and say this node has a “betweenness” of X or Y or Z.

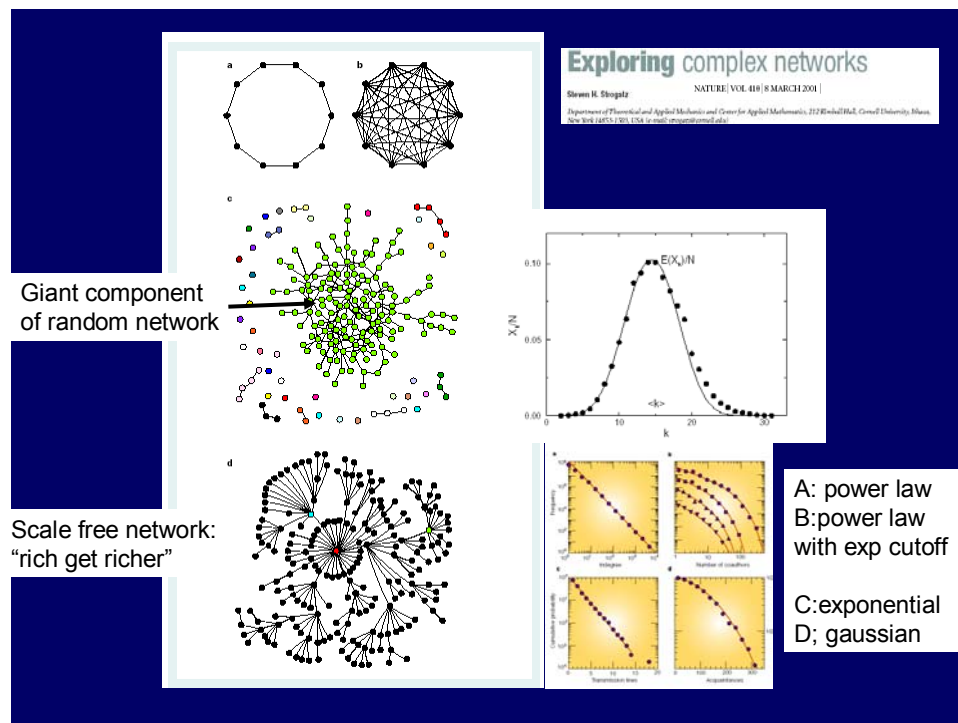


Figure 1

Figure 1 shows some examples of networks. This is from a great review by Steve Strogatz in *Nature* called “Exploring Complex Networks.”¹ On the top left, there is a very simple network, just a ring in which one node is connected to another. The one on the top right is a completely connected network; everybody knows everybody else. If you take out one of those nodes, the network won’t really fall apart, because everything else is tied in there, so the network doesn’t break up into something we call giant components.

The second image is what we call a random network. In a random network, you plunk down a bunch of dots and randomly connect the dots together. Steve Strogatz describes a giant component of a network this way: you plunk some buttons down on a table and start connecting the buttons together with pieces of string. Early in the process, you pull on the string and you pick up a couple of buttons. After a while, you put a number of strings together, and finally you are picking up a whole bundle of buttons, and that is what the giant component system is: the maximum connectivity of a subgraph.

Now for the scale-free network. If you look at the degree of some of those nodes, you will see it doesn’t follow a Poisson distribution, it follows what is called a “power law.” That means that a few nodes have many connections and many nodes have very few connections and there are real hubs in that system. Think about airlines and airports. Chicago has lots of connections, and there are very few Chicagos. Then of course there is East Islip Airport and Wings Airport and all the other little airports throughout the country and they have very few connections. That is where you get a power law distribution, which is not a random connection. The roadways and highway system are more of a random connection: a city develops here and a road is built, another city, another road is built, and they form a random network rather than a power law network. It turns out that networks that have a power law description will be completely different to attack and to defend. If that red hub in the middle of the bottom image in Figure 1 gets blown out, the network can fall apart into subcomponents. If you blow away the hub, then a guy on one end can never get to a guy on the other end. That is the idea of centrality in the network.

Let’s consider the distinction between failure and attack in random networks versus a scale-free network. What would happen in a random

¹ *Nature* 410, 8 March 2001, 268-276.

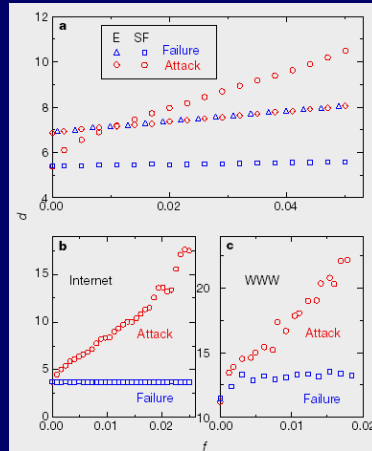
attack in a random network? If one of those nodes is blown apart, nothing much happens. Another one is blown apart, nothing happens. Finally after you blow enough away, the network falls apart. That's true whether the node is attacked or it fails, because there is no real structure in there, and the number of nodes you need to hit in a network before it falls apart is a characteristic of that network. Now consider hitting nodes in the scale free network at random and think of failure of a node as a random event. If I fail in a place that has few links, or a small degree, nothing much is going to happen to the network. So random failure in a scale-free network is going to look like failure in a random network.

But what about a deliberate attack on a node in a scale free network? If I could pick which node to attack, it is going to look quite different: the network falls apart. It turns out that that intuitive notion can be proven mathematically and that is what Barabási did.

Consider a random network versus a scale-free network, one in which we just threw down a bunch of nodes and connections versus one where we actually deliberately have a scale-free hub. What does the attack look like? This is exactly what I was telling you before. E stands for the exponential of the random network, F stands for the scale-free network, and blue is failure. Failure means a node, or link, (we've been talking mostly about nodes going out, but links go out as well with interesting consequences that differ from those of nodes being destroyed) goes out at random: for example, something happens to an airport randomly. An attack is a higher agent saying, "I'm going to go after one or two particular nodes." The difference is that in the exponential or the random network, you can see that as you increase the number of nodes that you take out, the distances actually do not change much at all. They are basically the same, whether you randomly take one out or you attack one, but since you are on a random network, it doesn't make that much of a difference.

The scale-free network, though, is quite different. A random failure on a scale-free network doesn't do much because you have enough nodes that if you take one out, it doesn't do much. But if you attack something, then the distance between the nodes really increases quite dramatically (Figure 2). That has been proven by looking at the internet and the World Wide Web. The fundamental difference here is that a randomly connected network behaves differently from a scale-free network with respect to an attack versus failure. That is really key, and I will show you why in a metabolic network of a serious human pathogen in just a moment.


Changes in the diameter of the network as a function of the fraction of the removed nodes



Comparison between the exponential (E) and scale-free (SF) network models. The blue symbols correspond to the diameter of the exponential (triangles) and the scale-free (squares) networks when a fraction of the nodes are removed randomly (error tolerance). Red symbols show the response of the exponential (diamonds) and the scale-free (circles) networks to attacks, when the most connected nodes are removed.

Figure 2

Degree distributions for three different types of networks



(a) scientific collaboration networks of biologists (circles) and physicists (squares);
 (b) a collaboration network of movie actors;
 (c) network of directors of Fortune 1000 companies.

From: Newman, Watts and Strogatz. Random graph models of social networks. PNAS 99, 2566, 2002

Figure 3

Figure 3 shows some more distributions of scale-free networks. The graph on the left shows the number of scientific collaborators. If you are from the American Institute of Physics (AIP), you probably already have this data in your database. How many collaborators do I have? It turns out that it is almost a scale-free network. There are many people who have one or two collaborators and a very few people who have lots of collaborators. In the physics community, there are millions of papers with thousands of authors on them. In the biological community, we have many fewer collaborators, but we still follow the same kind of general rules of collaborators. The network of movie actors again basically follows the scale-free network, and that is the “Six Degrees of Kevin Bacon.” It is a little different for networks of Fortune 500 or Fortune 1000 directors. You see, the x-axis is not logarithmic; it is a linear scale. The reason for this is still being worked out. It is important to understand these issues because how networks grow is a big question for the research community.

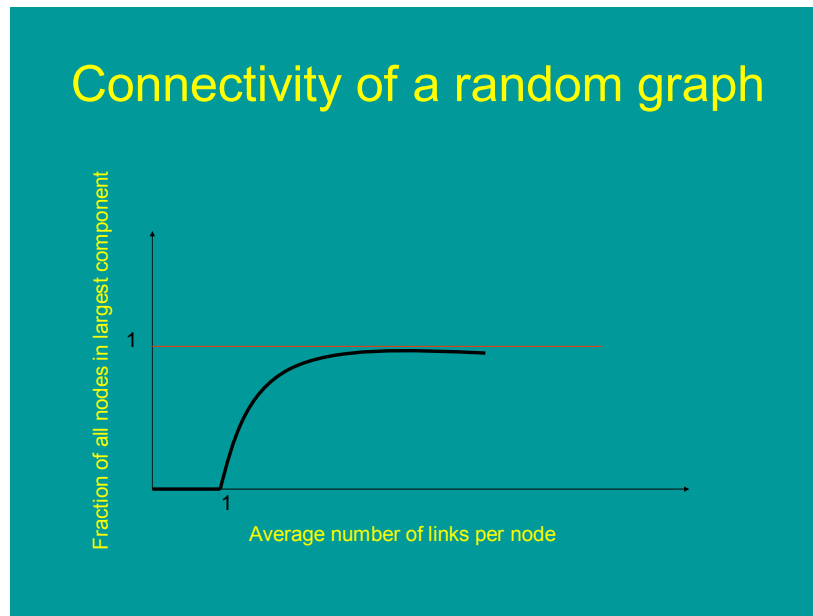


Figure 4

Figure 4, a nice graphic from Watts' book, shows a fraction of all nodes in the large component, and the fraction approaches one at a critical point. We can have more than one or two extra threads in the button system. Now why is this important? That becomes very important for dealing with the SARS (severe acute respiratory syndrome) outbreak: I am going to try to convince you that this graph can help explain the difference in the spread of SARS and Ebola. Now I have your attention!

The question is how many links do I have to add to get to that critical point? The image on the left side of Figure 5 shows a group of nodes hooked together and a few shortcuts. (This again comes from Strogatz's review article.) I can make a few shortcuts: to get from this node to that node over there, I can go in a line from one node to the next, or I can just jump across. Another example of a short cut is taking a plane from Hong Kong to Toronto. The graph on the right side shows the average path length, that is, how many jumps you have to make to get to another node, depending on the number of shortcuts. The more shortcuts, the quicker I get there. That is all that that actually means, and that curve is going to help explain SARS.

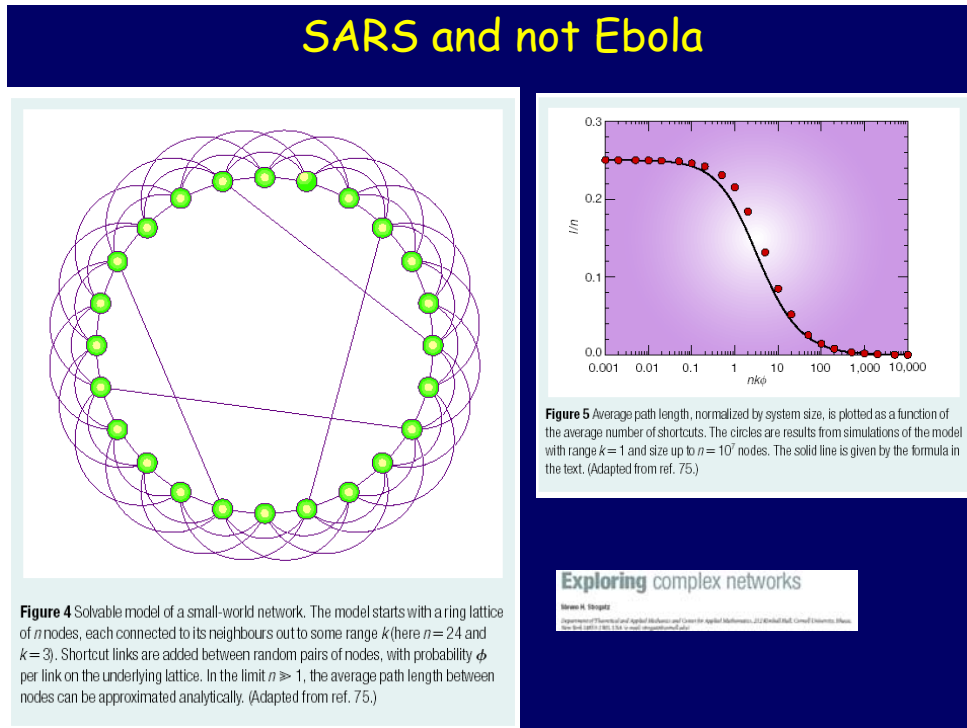


Figure 5

We have been talking about airlines and vaccines, but consider a germ like tuberculosis, which is what I work on. This germ has evolved to have this very complicated metabolic network that allows it to survive and cause devastating disease world-wide. In fact TB kills more people than any other bacterial disease. The interconnectedness between TB, AIDS, poverty and national development is a tremendously important problem, one that we are working on at Penn. If we could figure out how to attack one of the central parts of that metabolic network, we could take down the

bug. There are always mutations in the bug's DNA. Something is going on randomly in the bug that may inactivate it, but the bugs generally survive because these mutational events are basically random and it is not going to mutate itself out of existence.

Question: If you attack a scale-free network successfully, is it like a random failure in the hub?

Rubin: Exactly right. If you have a random failure in the *hub*, one chance out of a hundred thousand say, then you're in trouble. Then it's like an attack and you can't tell the difference. Not all attacks are bad, though. We just talked about an anti-bacterial or anti-cancer attack. If you can find the hub, then you are in good shape.

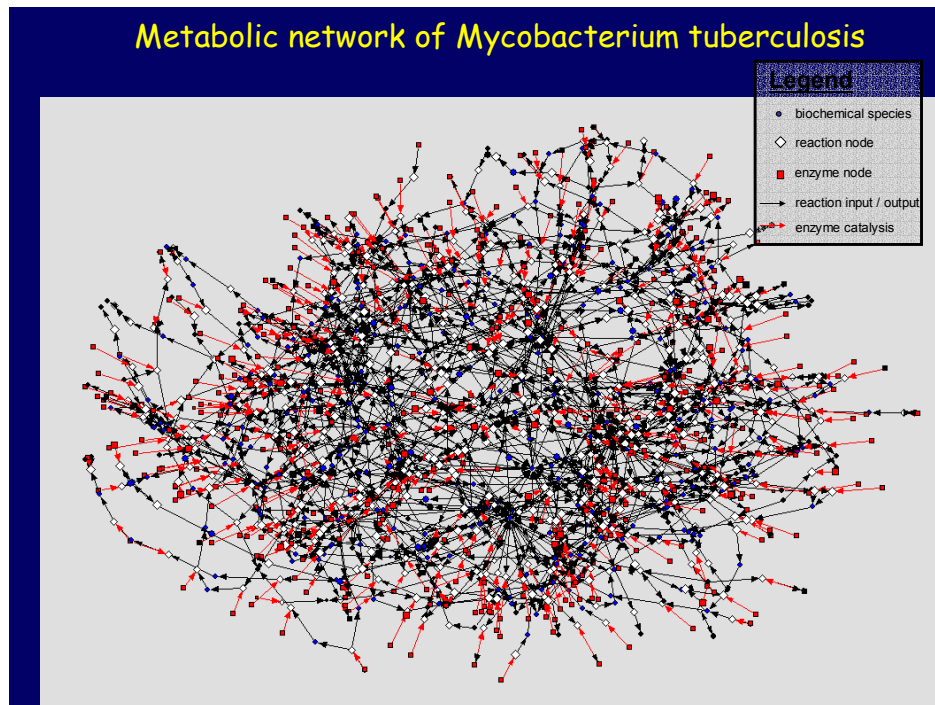


Figure 6

Figure 6 shows the entire metabolic network of *Mycobacterium tuberculosis*, as we put it together in my lab, and here I thank a very talented MD/PhD student in the lab, Marcin Imielinski. Every gene is known; the genome has been sequenced and there are about 4 million base pairs. We connected all the enzymes with their appropriate substrates and their reaction products and, to date, this is what the network looks like. As we

go along though, we are finding that even this isn't complete and we are continuously filling in gaps.

We are particularly interested in one sub-network that we call REL, which is the little red square on the upper right hand corner of Figure 7. That turns out to be one of the important hubs in the metabolic response in TB. We have attacked that gene and knocked it out with genetic tricks and the TB bug will grow for about two weeks and then die off, because that gene is really quite central to metabolism and many other things. I won't go through that; that's all published work. But it was a deliberate attack on the tuberculosis network. How we did this is another story, which I would be glad to tell another day. But it works. If we just waited for nature to randomly attack and hit that one, we might wait forever. And in fact, it would not survive as a pathogen, because it would just die out. So if a bug is randomly attacked in its central hub, it is not going to survive. What we have to do, as molecular biologists and physicians, is find those hubs, for example in cancer networks or bacterial or viral networks, and design a way to take that hub down – these techniques and ideas will be used in future drug discovery and in designing treatment programs.

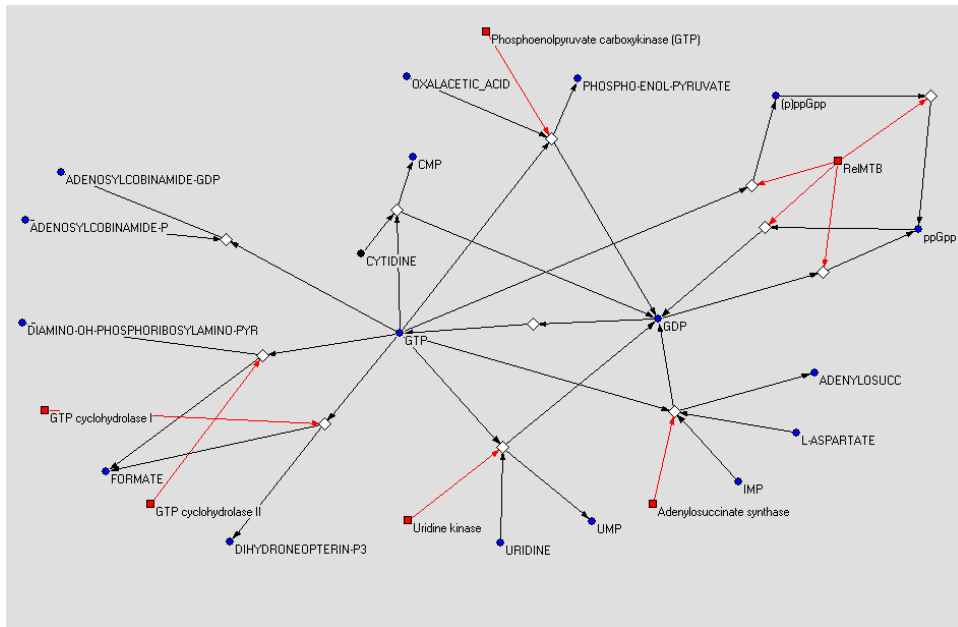


Figure 7

So far I talked about *static* descriptions of networks. The really hard problem lies in the *dynamics* of networks. We can perhaps knock

down a node or link for a short while, and the network will fail for a while, but if you take away that tweak, the network will come back. So the time-series and the dynamics and the kinetics of networks are very important and interesting, and the reason that that is so hard to do, at least in the stuff that I work on, is that we don't know the kinetic constants of many of these reactions yet. We spend a lifetime working on one single enzyme! That's why we need a systems biology approach, collecting and curating all the data from the entire community of scientists working on a problem.

Question: I understand the node concept in a transportation network and the internet. What are the nodes in a bacterium?

Rubin: Reactants, or metabolites in an enzymatic or genetic pathway can be the nodes. The physical world is connected and physical things move through those connections. Cars move across roads and baggage moves through airports and metabolites move through reaction pathways. Our work is on the enzymes that connect one substrate to another. The boxes and triangles in Figures 7 and 8 represent enzymes and reactants and reaction products. That is essentially what makes you different from a stone. You can catalyze reactions that a stone can't. This operates at many different temporal and spatial levels. You have one network here and a regulatory network here and they work on different time scales. Sometimes they work in different places as well. We understand a lot about a regulatory network on one level, both in time and space, but when you put on top of that a controlling network, then it becomes very, very difficult because different time scales are involved. Metabolic reactions take place at milliseconds and nanoseconds while the regulatory network that is being turned on and off goes in minutes and hours. You have to be able to compose these two networks somehow, and mathematically that is very hard.

Think of it this way: you are an engineer building a road from Washington to Philadelphia. You get your guys out there and tell them to get going as fast as they can. You provide the raw materials and the gang gets moving. Then all of a sudden the lawyers call you and tell you to stop because you can't go across some jurisdictional boundary without first obtaining the legal right of way. What do you do to optimize your work? Do you tell your crew to slow down, but continue to work hoping that getting the job done will coincide with getting the right legal okay? At least your guys are working and getting a pay check. Or do you tell them to stop and go home – no more immediate waste but what you've built may start to decay without further work? Or do you tell them to speed up, get to the boundary and then we'll worry about it, maybe force the hand of the legal or political system? All these are a priori reasonable solutions to the

or political system? All these are a priori reasonable solutions to the optimization problem, but only one may lead to the survival of the project or your career as the lead engineer. The same problem exists for the bug – how to get along with different regulatory influences.

Now let's consider other catastrophic failures. One option is to go through all these catastrophic failures and try to figure out why they happened, for example, Three Mile Island, petrochemical plants, aircraft and marine accidents, the space shuttle, and the power grid failure of 1996, in which eleven western states and two Canadian provinces lost power. You can read about them in two recent books, *Minding the Machines*, by my colleague at Penn, William M. Evan, and *Normal Accidents: Living with high-risk technologies* by Charles Perrow from Yale. They go through a whole bunch of different catastrophes; it's not fun reading. The power-grid failure is discussed in Watts' book as well.

Let's pick a big one, the East Asian economic crisis of 1997. I just finished reading Joseph Stiglitz's book *Globalization and its Discontents*, and whether you agree or not, it is a terrific book!

Consider the East Asian economic crisis of 1997
Here are some statistics: People living in poverty increased by 100 million in a time when the total world income increased by 2.5 per cent annually Economic crises in the last fifty years have been more frequent and deeper—almost 10 countries have faced crises Every major emerging market that has liberalized its capital market has had at least one crisis. Claim: IMF adopts the Washington consensus of free markets and the right policies for developing countries. IMF funds and programs have not stabilized the crises but made them worse. <i>Globalization and its Discontents</i> , Joseph E. Stiglitz

Figure 8

Figure 8 shows some of Stiglitz's statistics. You can see where he is going. If I read him right, he says that the IMF adopted the Washington consensus of free markets and correct and right policies for developing countries, but he says that IMF funds and programs have not stabilized the crises and made them worse. So why is this a complex network of interdependent security? He gives the answer in a paragraph:

“When the Thai baht collapsed on July 2, 1997, no one knew that this was the beginning of the

greatest economic crisis since the Great Depression — one that would spread from Asia to Russia and Latin America and threaten the entire world.”

“I believe that capital account liberalization was the single most important factor leading to the crisis.”

... “By continuing to advocate contractionary policies the IMF exacerbated the contagion, the spread of downturn from one country to the next. As each country weakened, it reduced its imports from its neighbors, thereby pulling its neighbors down.”

“As the region imploded, the declining demand for oil and other commodities led to the collapse of commodity prices, which wrought havoc in other countries, thousands of miles away, whose economies depended on the export of those commodities.”

That is a classic definition of an attack on a hub in a complex network.

The argument is basically this: rapid trade liberalization and opening up to imported products that compete with local produced goods *should* enhance a country’s income by moving resources from less productive uses to something the country can actually do. But this has profound consequences on social and economic structures: jobs were not created and the tight monetary policies actually did not bring money into the country and kept money out. So it’s an example of a network failure, with tremendous economic impact.

SARS is exactly the same thing. Just pick up the paper and you’ll see it may have a huge economic impact and may be the new East Asian crisis. I will make the connection now. Just put in a few shortcuts and you have a collapse of the network. Bill Marsh at the New York Times, had this graphic (Figure 9) about two months ago, which I am using with his permission. Here is that shortcut. It starts when a Chinese medical professor who had been seeing SARS patients in Guangdong Province becomes infected, and he travels to a hotel and infects everyone else. Then these people travel elsewhere. Had that doctor not gotten on that airplane and gone to the Metropole Hotel, SARS may not be with us today. In effect, it would stay localized in the community and die out in the community. I will come back to this.

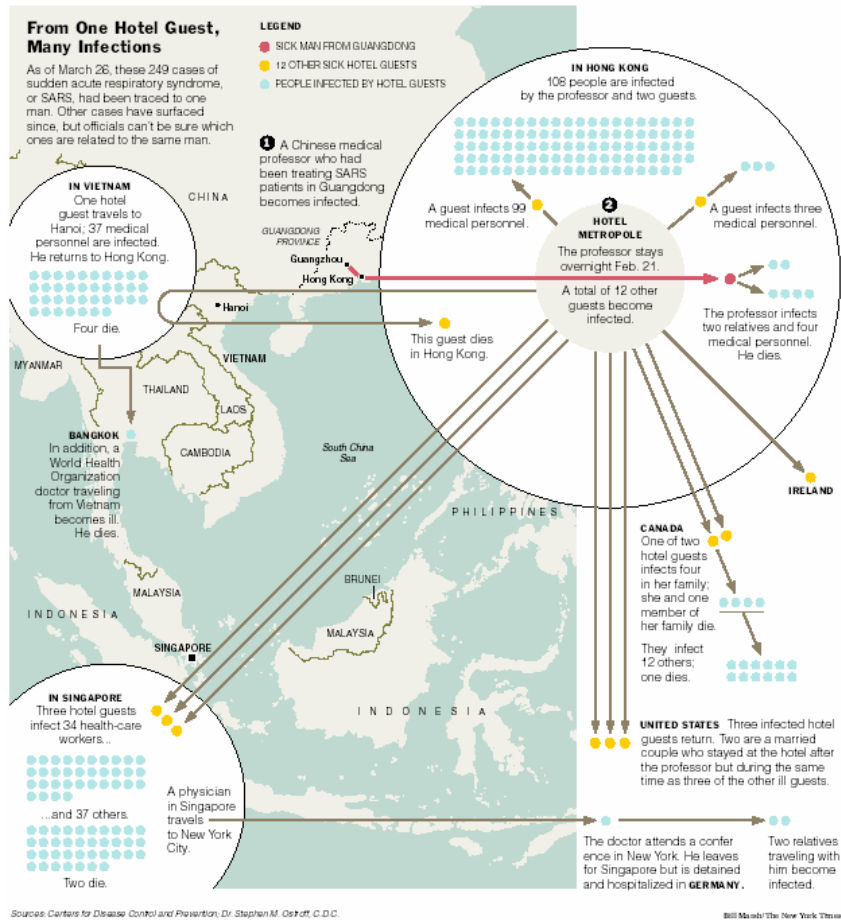


Figure 9

So this is SARS and not Ebola. Why is it not Ebola? Because you don't get on a plane with Ebola. By the time you're infectious with Ebola, you're basically dead; you're hemorrhaging massively and it is direct body fluid contact that spreads Ebola and that is why it always contained, that's why luckily we haven't seen Ebola in this country yet. With SARS, you're infectious before you know it, and so you can get on a plane and infect people. It still could happen – it is more uncommon, but it still could happen. I am convinced that this is the explanation. People ask, "How did SARS spread?" It spread because it's a small world, because of shortcuts. It reached that critical point right there and we are still seeing the consequences. It is only when you break those shortcuts or you put in infection control measures that you can contain the outbreak. It may be that we are getting there now. It will spread all over the world, and I just heard that

there are five or six more countries that are reporting SARS cases. Vietnam actually had its last case two weeks ago. As of April 23, there were 3,098 cases and rising.

Now for a counter-example: with the Black Death in the 14th century, there weren't that many shortcuts. You could not jump on a plane and spread the plague, but in fact, it did spread in waves. The Black Death spread from the Black Sea and then it moved in waves up through Western Europe until it destroyed up to a third of the population. There was a nascent network; the evolution of the shipping industry at that time allowed ships to travel around that part of the world even during the wintertime. So there was a technology that advanced, much like airplanes, but with ships, so that the shipping network connected the Mediterranean world with the rest of Europe. Year-round shipping, the introduction of black rats into Europe, that whole network constellation allowed the spread of bubonic plague. Other things that were going on as well: deforestation and people starting to live in concentrated groups like manor houses and towns that weren't so scattered around, all of which is described in Norman Cantor's *In the Wake of the Plague*. There was a tremendous impact on society as well, with great social upheaval and the rise of movements like the Flagellants. The Jews were blamed for the plague, of course, and were pushed into Poland and the Pale. There were all sorts of impacts on literature, the arts and the economy as well.

The question is can anything like that happen today? Certainly SARS is not going to wipe out the world; we will be able to contain it. But there is a disease that is devastating a large part of the world, and that is HIV/AIDS. If you look at five countries, Russia, China, Ethiopia, Nigeria and India, there is already a tremendous economic and social impact of the disease. By 2010, these countries will suffer decreased productivity and profitability, decreased foreign investment, a drop in the GDP by 20%, and Nigeria could be impoverished in the next couple of decades. Kenya's GDP is expected to decline by 15% by 2005 and Tanzania's by 20% by 2010. We expect a decrease in human life expectancy by as much as thirty years, and a fourth of the population will die over the next decade, leaving a huge orphan cohort. The impact on the social structures of these countries will be almost beyond belief: 42 million children in twenty-seven countries will lose one or both parents to AIDS by 2010. That is an astonishing number. 23 million individuals are now infected with HIV, which will rise to 75 million in 2010; one-third of all military conscripts in these countries are unfit for service due to HIV, hepatitis and drug use. This data comes from a National Intelligence Council report in 2002. They actually see this as a mili-

tary threat as well, since you can't sustain a functioning military when both the young men and women in the military have AIDS.

Economic consequences of naturally occurring diseases	
Avian flu, Hong Kong 1997	\$100 million in lost poultry production, air travel off by 22%
Avian influenza in the USA	\$63 million
BSE, United Kingdom 1995	\$9-14 billion
Cholera, Peru 1991	\$775 million lost in ban on seafood exports
Plague, India 1994	\$2 billion in losses to Indian economy – half a million people flee – aviation, tourism shut down
Hypothetical smallpox attack	\$117 billion/ week

Figure 10

But the disease doesn't have to be AIDS: any naturally occurring infectious disease will do something like this. Figure 10 shows some of the results of recent epidemic diseases.

What are the Threats?	
Biological Agents (Category A)	Anthrax, Botulinum Toxin, Plague (<i>Yersinia Pestis</i>), Smallpox (<i>Variola major</i>), Tularemia (<i>Francisella tularensis</i>), Viral Hemorrhagic Viruses
Chemical Agents	<u>Nerve Agents</u> : Tabun (GA), Sarun (GB), Soman (GD), VX <u>Blistering Agents</u> : Sulphur, Nitrogen Mustard (HD) (HN), Lewisite (L), Phosgene Gas, (CX) <u>Choking Agents</u> : Phosgene (CG), Chlorine, Chloropicrin, (PS) <u>Other</u> : Hydrogen Cyanide, Ricin
Conventional/ Unconventional	<u>Nuclear Weapons</u> : Conventional, Suitcase, Dispersion (Dirty) <u>Conventional Explosives</u> : Suicide Bombing, Vehicular <u>Sniper Attacks</u> <u>Transportation Sabotage</u> : Hijacking, Rail lines <u>Cyber Attacks</u>

Figure 11

Figure 11 shows some of the other threats that the CDC and World Health Organization have put together. This is what we have to worry about, in a nutshell. It is a daunting list, because every one of these has network consequences that include psychological issues, commercial issues, economic issues, and so on.

An article in yesterday's *New York Times* showed an example of a very interesting network that has become very popular: the control of

global resources as a way to influence a network of individuals or cultures. It is the protection of that network that we can consider a component of post-conflict reconstruction. Since this is the George Marshall Institute, we should pay homage to the man who virtually invented post-conflict reconstruction. This is another example of putting together complex networks and making sure that they are robust and secure. My colleagues at Penn, Robert Giegengack, Thomas Naff, Christopher Williams, Chris Beals and I put together a proposal to work on the river systems and health consequences in post-conflict reconstruction in Iraq.

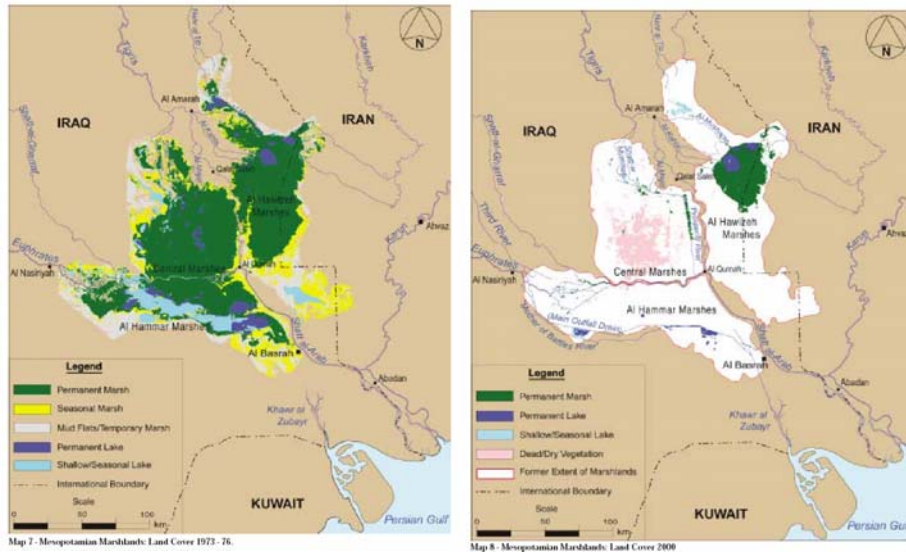


Figure 12

The marshlands in Iraq in the southeast portion of the country, on the southern bank of the Tigris River, were an amazing place; some people even think it was the original Garden of Eden. People constructed floating huts out of reeds and built beautiful villages with amazing animals and birds and fish. They lived like that for millennia and were Shia Muslims, who obviously were in opposition to the Sunni Muslims of the Ba'ath Party. Saddam Hussein said something like, "We don't need these guys here; they're just a pain in my southern flank." He started to drain the marshes. Half million to a million Shia Arabs then migrated into Iran, and there are maybe 20,000 left. The marshes are now almost completely dried up.

80% of the Tigris arises in the mountains of Anatolian Turkey and about 60% of the Euphrates arises in Turkey as well. Turkey could turn off the water to Syria and Iraq tomorrow by shutting the big Atatürk Dam on

the upper Euphrates. It would take a couple of weeks to get all the dams in place again, but they really could turn the water off. Saddam Hussein actually turned off the water to the dark green area in Figure 12, causing the destruction of the marshes: the picture on the left shows the original marshland in 1977, and what remains is on the right, about 10% of the original marsh.

It's a huge humanitarian catastrophe. The Marsh Arabs' way of life, their culture and society, their health and well-being have been negatively impacted by draining the marshes. Why did he want to drain the marshes? Political control of the population, oil In fact, after he drained the marshes, he set fire to the reeds, making the area totally uninhabitable. But he has also taken most of the water and started irrigating northern parts of Iraq, which has completely changed the environment.

How do we reconstruct the marshes? This is a major problem which involves ecology, the environment, hydrology, geology, and medical, social, economic, and cultural issues. Bringing the Shia Muslim population back from Iran into the marshes and dealing with their health conditions are major problems. That is a project that my colleagues at Penn are working on and it is going to be difficult. It is not clear that those marshes can actually be reclaimed and that you can rebuild this society there.

So again, this is an example of the interconnectedness of network analysis. We can quantify these things, we can measure them, and we can predict and plan for various attacks and various failures. It will take the creativity of everyone, like the people sitting in this room, who know about security, who know about medicine, economics, politics and law, to sit down together and figure out the structures of the networks. We certainly now know that hierarchies of networks behave in unusual and as of now, unpredictable ways, so it takes these kinds of colloquia to really have a constructive dialogue. Thank you for your attention to this discussion.

* * *

Questions and Answers

Question: I want to go back to your theory. I don't know if you answered this question in your published studies. Most of your work and most of the examples provided in your data or other data come from Third World countries where there are governments that have authoritarian rule. What about the dynamics of democracy in all this? When you look at injury data from nuclear, biological and chemical attacks, that data mostly comes from Rus-

sia and similar countries with those kinds of governments. What about the dynamics of an educated democracy? Can we respond appropriately on an individual level? When we get to that, in a network, how do you build your flexibility in your response, in your network?

Rubin: It's an interesting theory that maybe democracy and liberal society make you more resistant and resilient to attack. I think that one can propose the opposite as well, that the more open the society is, the less resilient and the less resistant it is to attack. And we can ask whether controlling that threat is worth all of the other problems that are cropping up around the possible abridgement of constitutional rights and free speech. It is an area that will require detailed research and it is a very interesting and important question.

Question: The al Qaeda network has different cells functioning as independent nodes and they really don't have a network, so we could not on the spot totally kill it. If you actually have a network and the hub is hit, can't any of those nodes break off and become independent?

Rubin: That's a good question. I am not an expert in that area and I don't really know the answer, but I am willing to speculate. The reason that al Qaeda was so successful was that it probably didn't depend solely on bin Laden, and if you take out that hub, it is more like a random network; the nodes can act independently. Why didn't we have terrorism happening in the United States during the war on Iraq? Possibly because there were not even ideological links to Iraq. Possibly it was a question of timing.

Question: You brought up the subject of intellectual freedom for scientists, and the problems with that. The scientists in Australia two years ago wrote about interleukin-4 and how to make smallpox resistant to our vaccine. Two Army pathologists today published the code for Spanish influenza. Johns Hopkins has done a big study on that. I understand your comments from your perspective in the scientific community, but for me, from the national security community, it is really frightening what you can find. I downloaded the genetic code of the Spanish flu off my daughter's computer! It's insanity to have to have that out there. How do you stop it?

Rubin: I will give you a worse example and some more bullets for your gun. The Russians clearly knew how to engineer anthrax as early as about 1994. We had always thought the Russian biological community was far behind us in technology. After all, they were remnants of Lysenko days. They were good mathematicians, good engineers, they know how to play chess –

that's great, but they still don't know how to do molecular biology. Then this paper comes out 1994. The two Russian scientists were asked how they knew how to do this. Do you know what their answer was? "We read the papers in the published literature." The Russians signed the biological warfare agreement and they went merrily along continuing to do biological warfare development, and not only merrily, they ramped it up. Did our guys ever think that what they published would be used in some nefarious way? I don't think so, but where do you draw the line? My sense is that we can't hide all the information that we have. It's in the open community so that people will be able to work on it and think of countervailing measures. But certainly, my colleagues in the cyber-community and the code community aren't allowed to publish some of their work. My sense is we need to have intellectual freedom and that's the only way we will stay ahead.

Question: Can you publish a summary instead of all the details, though? I was just looking at Johns Hopkins and the idea that there are certain people who can give input. It doesn't have to be available to everyone.

Rubin: That's a compromise. Again, we have to put our information in the public domain to have it judged and see if it's right or wrong. Let me give you a small example about the way I got going as a young assistant professor. A researcher published a sequence of a gene that I was working on. His was off by twenty base pairs. Talk about catastrophic events! At the same time, a researcher in Germany was trying to solve a crystal structure of the protein that the gene encoded, but he couldn't solve it based this other guy's published sequence because the amino acid sequence was wrong. We published our sequence and Dr. Bode was able to solve his crystal structure. He invited me over, actually we escaped from Slovenia together the day after their independence was declared, but as Kipling says—that is another story. Dr. Bode and I collaborated, one thing led to another and that's basically how my academic career got started. Just by twenty base pairs. You might say the devil is in the details, but that happened to be very important. Who knows?

Question: There was an analogous problem before the Second World War. There were all these studies in radioactivity and physics and there was a serious problem for *Physical Review* to be published at this time. The community actually exercised self-censorship.

Rubin: It is said that it is hard to do some of the physics bad-guy stuff; even if you know the physics, it's hard to get that bomb-making radioactivity

production going in your backyard. It's hard to do. Whereas biological stuff is not that hard. You can buy all of the agents you need from kit companies now, where high school teachers buy supplies for their students. So there is an issue of self-censorship on the physical side and it may need to be there. But with biological science, the argument is that we need to have to have more censorship because it's so easy to do. I don't agree with that argument.

* * *

RECENT WASHINGTON ROUNDTABLES ON SCIENCE AND PUBLIC POLICY

James Oberg – *Toward a Theory of Space Power: Defining Principles of U.S. Space Policy* (May 2003)

Sallie Baliunas and Willie Soon – *Extreme Weather Events: Examining Causes and Responses* (March 2003)

Randall Correll – *National Security Implications of the Asteroid Threat* (February 2003)

David Trachtenberg – *Weapons of Mass Destruction and Terrorism* (November 2002)

Henry Cooper – *Defending American from Offshore Missile Attack* (October 2002)

Baker Spring – *Missile Defense in a World Without the ABM Treaty* (June 2002)

Gregory Canavan -- *Update on Missile Defense Technology* (May 2002)

Dorothy Denning – *Is Cyber Terrorism Coming?* (May 2002)

Jesse Ausubel -- *Does Energy Policy Matter?* (April 2002)

Bruce Ames – *The Causes and Prevention of Cancer: Do Federal Regulations Help?* (March 2002)

James Schlesinger – *Responding to National Security Threats in the Post 9/11 World* (February 2002)

The Marshall Institute – Science for Better Public Policy