Event ID: 1918140
Event Started: 3/15/2012 11:48:03 AM ET

3/15/12 [Captioner present. Please standby for real time.]

Hello. I'd like to welcome everybody to the 10th in our monthly series about WATCH security. In the past few years I've heard lots and lots of talks about security in our community. And there be also a lot of bickering about what the definition of science is. From my point of view, the basic idea is 2,000 years ago the model of research was you had smart people making observations and thinking a lot and making pronouncements. He pronounced men had more teeth that women did and heavy objects fall faster than light objects. And eventually people realized this just wasn't a proper way to gather knowledge. And instead of having smart people, you should make it clear why you believe what you believe. And the importance is not who said it but why they said it. And we can question the reasons why and make guesses that way. Unfortunately we all realize that doing science is more difficult than we would like. So our speaker Tom Longstaff is going to talk about the barriers for doing that kind of research. Tom is the chair of computer science information assurance and engineering programs at the John Hopkins University and technical director of the systems behavioral group at the MSA. So I'd like to turn the talk over to him.

Thank you very much. Can everybody hear me okay. ? Is that better?

All right. Normally people don't need an MIC to hear me but that might be tough for the folks on the line. I can project pretty well but not to California. Thank you very much for having me here. This is quite an honor to be able to be part of this series and be able to talk to you about one or two of my soap boxes. Let me set a little bit of context for what you are about to hear. As Sam said, there have been dozens to hundreds of talks on science in security. I think I've personally about to about 20 separate talks on science and security over the past decade. All the talks have the same kind of idea. If we only did science and information security we would be so much better off. We would have natural laws that would be functional. And send people to the virtual moon of the information strike that os fear. And -- STRATOSPHERE which usually isn't very scientific. So for the last two to three years I've been looking hard at why can't we actually really engage in a more scientific way in the information community? What is really keeping us from doing that? What is the problem if it's something that's universally agreed that would be the same? Why can't we actually do it? I will say right now in the interest of my definitions of what I'm talking about, this is not a talk about a scientific study. I have not done a scientific study of why we haven't done science. This is a series of observations and interpretations on my part. If any of you disagree with my interpretations, I encourage you to form a hypothesis and form a study. That would be a useful thing to do. But consider this talk highly observational. The reason why I want to do this talk here is [Inaudible]. I've identified a number of relatively small things that we BOETHS within the U.S. government and within the field could do that overtime would have a large impact on this problem that we have in the area of science. It's not just a whining session. I don't expect to sit here and whine. I do expect to make recommendations that each one of us can do to help change this over the period of the next decade. And make this a kind of scientific endeavor we all hope it will be at some point. And the slides work.

As we have mentioned there have been numerous talks, studies, papers, books, things written all over the place written to cyber security. We have admired this problem for a long time. We have written again and again. These slides mostly came from a talk I did in 2010. And since 2010 since I put this slide together, we've had a few more things that have sort of come out. Most of you have seen the Jason’s report. The Jason's report was DOD funded study that really began to look into is there a science to be had here in security? And the conclusions of the report are yes, there is science within the study that we're doing. There is a potential for experiments of a scientific nature. However, the report says there are no fundamental natural laws to be discovered which I thought was an interesting conclusion. But that is one of the conclusions. Unfortunately, there is very little in the Jason's report that says how do we get there from here? So again, doesn't really tell us where to go. More recently, I've seen this great paper supporting experimental computer science which is an internal report. Again, this is very typical of reports that we see that call for science and conclude by saying we need another test bed. And again, this is not exactly where I think the hard part of the problem is. If the problem were just building a test bed, we would have done that. And we would be a scientist at this point if that was all that was required. But clearly, we need something more than that in terms of where we go. We're stuck in admiring the problem, stuck in analogies and metaphors in terms of trying to move forward. I've seen numerous comparisons, chemistry to physics to mathematics, to medicine to on and on and on. Lots of analogies, lots of metaphors. Nothing very actionable in terms of where we're coming from here.

So my observation through the last three years is that there are three main barriers that are keeping us from getting there that I want to talk about. And what I'm going to spend the most time on is time. That's where the majority of the next 40 minutes of the talk are going to be is in this particular area. There are lots and lots of implications on this we have to address. This also leads in the standards of peer review and expectations in terms of where we're going. The scope of where I'm talking about, the central idea here is that if we can overcome these three barriers, we have a far better shot at actually achieving information security research. And actually getting there from here. These are the most important barriers I've found. You may imagine some other ones I won't talk about. I've heard many other suggestions. These are the three that have seemed to be the most in the way of kind of what we're doing. That's the central idea. Overcome these barriers and we can actually get there from here.

Now, Sam said the definition of science. I'm not going to give you a formal definition. I was a trained physicist. I came to computer science late. I moved from the world of light to the world of darkness into this area. I had a strong bias toward experimental science. A lot of my observations are tainted with that particular view. I have not been a theoretical physicist or theoretical computer scientist. However, my observation has said there are many publications in computer science and information security that have some of the same problems I'm going to talk about. Although, I won't address them directly, you might see yourself in some of those areas if you are more in the theory kind of world.

So let's talk about time. Bottom line, experimental science takes time. It takes a lot more time than we're willing to stomach. We have no stomach for usually doing the hard work that's necessary to do experimental science. Why don't we have that stomach? Theres a whole lot of reasons. I'm going to go into the whys in a minute. Anecdote first. Before I started this new position where I'm working in DOD, there was a time that I was leading research at a major organization. And I was actually sitting down with a room full of people who were actually PHD level researchers in the field of information security. And I was asking this room basically let's talk about the current project and the next things that we really ought to try to study. More scientifically, what we're doing. And I got this as a real answer. How will we know what to study until a call for paper comes out? How will we know what kind of research we can public until we see the call for papers? And this completely blew my mind. You are suggesting we start research when the call for paper comes out and complete it by the due date. Right. That's the way it works. Perfectly, could not see the problem with this. And I thought maybe this is just that one case. But no, I looked around and I've been on a lot of programs and done a lot of various works in conferences and different kinds of publications and I see this over and over again. People begin to write when the call for paper comes out simultaneous to doing the research it to put into the paper that it's an information security paper, I can finish it by the due date. And I can quickly have a graduate student write a program in time to actually submit the paper.

Now, you might think this is kind of an ought liar but it -- out liar. A lot of us don't think beyond the next existing contract that we might have. We don't think in terms of building on the last contract. We think about we think of a new idea, I'm going to submit it and get the money and in an 18 month period of time. Publish my results and be done and move on. I hear laughing. You've all seen this. You can't do good experimental science in this kind of approach. It takes time. Why?

All right. I hate having to put up this slide. I hate it. Because there shouldn't be a single person on the call or in the room that would look at this slide and say that's what scientific method is. But unfortunately, what I see frequently, my observation is then people trained in information security do not understand scientific method. They do not understand what it means to be scientific. They have been trained as programmers first. Software engineers second, large system integrators third. Creating a large system as part of PHD and launched into the academic strike that -- STRA TSHGSTOSPHERE. And I have seen seasoned researches attempt it on the spot. This is what I'm trying to test. Doing the results, do the conclusion, writing the paper is the least time consuming step in this process. It's the least time consuming step. Even a journal paper that takes a long time to write is less time consuming than doing this work correctly.

So if we have to return to basics and go back to what it means back into scientific method, what can I say about why we're in this situation where we can't form good experiments? And some of it comes down to this idea that the experimental side of computer science rarely, I won't say never. Rarely confirms with the theoretical side. Rarely. I have seen it in classical computer science. So cueing theories. We had great theoretical science that were debunk by experimental methodology. It was a fabulous result and caused a scientist to go back and redue what was really going on. It was a real advance in the field. I can think of very few that have followed the same idea. Decree tolly -- CRETOLOGY. For the most part we don't do it very much in this kind of world.

We certainly don't talk about causal relationships in this kind of an area. The experiments we do put together are largely demonstrations of some innovation we are creating as opposed to verification of a theoretical causal relationship in the environment. You look at any other sort of experimental field. Most of the experiments you see are about understanding the causal world we live in. We believe we're in a world of cause and affect. That's the fundamental thing in science. We can relate that and that will help us predict what the next action ought to be and move forward. Start with observation and mental exercises. But that's kind of the root cause of all that we're trying to do.

So we have to go back and identifying the experimental need. If it's about causal relationships, why do an experiment? I rarely hear this out there. Why do an experiment? We have built something, we must do something. I have to do an Experiment for -- experiment for my class. But rarely I hear why do an experiment? What is the point? How can we separate out this idea of observation from experimentation. We see lots of really cool stuff. From the cool stuff that we see, we should say, you know, I think A causes B. I think there's something here that relates. Whenever I see this one, I always see this one. Maybe I should do an experiment to find out the limitations between that causality relationship. Maybe I should sort of figure out and test these conditions. This would be a really foundational and startling thought to most people in information security if they would want to go in this particular direction. They would hypothesize about a causal relationship and go out and test it. It's an absolute natural way of thinking if you come from the physical science. This is the way you are taught to think. It's the way you are taught to critique other people's work. The way you are taught to actually move down and make progress is in this kind of way. But we have not done it in information security in especially not in sort of modern times. Where we have the internet as our primary area of work.

Now, I talked about formulating hypothesis already. And doing this on the spot. There are lots of ways to enable an experiment. Probably the best way I know of is to actually come up with a very clearly defined hypothesis you are trying to test. You can do demonstrations and do experiment that end in demonstration. You can do other sorts of things that drive an experiment. But you really want to have a clearly stated hypothesis that can be wrong. Because in fact, all you can really do is show a hypothesis as wrong. Right? Can you absolutely show the hypothesis is true? No. You can confirm the conditions you are testing but you cannot prove the hypothesis correctly. You can gather a lot of evidence for it. We go back to the Newton's days. You can say the original gravity of Newton was a great hypothesis and held true for a long time but wasn't necessarily the truth. And now we know Einstein may not necessarily be the truth. We don't look for truth. We look for confirm hypothesis. Things we can support overtime. However, refeuding hypothesis should be something we do all the time.

Has anybody here read a paper that's refeuded hypothesis? Almost doesn't happen. Cool. There is a conference coming up that talks more about this. That we actually are going to have a bigger body of knowledge in terms of trying to find things that didn't work, hypothesis that fails. That's a really important element that we just don't capture.

After you've got your hypothesis, you design your experiment. Here's the problem with this particular aspect. Somewhere along the line in the field, we equated designed the experiment with built the test bed. To fund larger and bigger and more complete test beds. That will design our experiments we're going to use. Provide the environment which all of our experiments will run. We have a whole variety of things. All of these test beds out there. That have been used as almost the substitute for designing an experimental environment. What's the problem with that? These test beds are general and people do good work on them. Why is that a problem? When you are handed a test bed before you've designed your experiment, it has already fixed several things that you can and cannot test. Requires that you have extremely detailed -- something that requires really details timing. Can you run that on a national test bed? No. Why not? Because it actually requires special purpose test equipment to accurately define what it is you are testing and gather the information to support or refute the hypothesis. Test beds don't do that for you. The other thing they tend to do is sort of make uniform the data that's available for tests. We have automatic generation of this information that's out there. You can run any software you want, test any tools you want to do but here's the data. We have all used some of these test beds that have user models that create user background traffic and data. And they are reasonably generic. They are supposed to represent the environment. Representing the environment is actually a function of designing the experiment of the -- experiment. It's not about one size fits all. About defining things appropriate. So this is a real problem in how we actually decide we're going to run experiments. I'll tell you right now that if one of the things you worry about is scale, there are very few places in the country or the world that can run something to scale, if that's an important part of your experiment. Are national test beds bad? No. But for them to drive the only experimental place that we support is probably not in the best interest of the field. Certainly, not in terms of recreating how we design experiments.

After you've designed the experiment you configure the experiment. I've spent a lot of time calibrating devices and equipment. And understanding exactly what I'm going to measure in the various kind of experiments. That was a common part of configuring an experiment. And making things work. In information security, we equate configuring the environment with creating the innovation or the technology. This is kind of where we do it largely. We create things and configure it so this is where we create the defensive tools that are going to run in this environment. This is where we create the measurement pieces we're going to do to extract information from the environment that we do. Sort of bundle all this stuff up together rather than say here's the experiment we want to run. Here's the thing we're trying to learn. Here's the information and the data that we need to run that and now we can figure it so it all kind of works in the way so it does the appropriate extraction.

In here, data control and characterization is really key. And it is far more important than most of the experimental designs give it credit for. We bury it in words like ground truth. We're not going to use ground truth. Is ground truth important to actually get the result you need from the test? If it is, then saying I'm not going to worry about ground truth is too hard, means you haven't designed your experiment appropriately. You can't do that. Got to make it happen.

Data characterization really the key to this sort of thing. Understanding extraneous elements. Here's where the test bed problem comes into play.

If you need to use something like a national test bed, one of the things you have to be very careful about is what else do you get in an experimental environment like that that you have to explain or count as another variable or try to control in some other way. It makes it really hard to sort of say that you have this environment where you've got way too many variables to actually measure what I'm trying to do. So I run a class in intrusion detection. One of the things I do is make my students design an experiment using deter. We have to design an experiment and they have to tell me what extraneous elements exist that might invalidate the results of their experiment. And that turns out to be one of the hardest aspects of the entire project. And actually, very few students get it entirely right in terms of what's happening. But it's one of the things I really try to work with the students on. Because I think they are going to end up in environments like this when they do their own experiments later. They have to understand how to do that in terms of trying to make that work.

Execution. So here's another area where training really fails us in the information security world. On almost every case that I can think about with only very few exceptions, people who talk about running an experiment and information security do it pretty ad hoc. How do I verify they've done it ad hoc and what's my observation that allows me to make that statement? One, I ask for the lab notebook of the detailed note of how they built and executed the experiment. Two, I ask for the detailed instructions that was required for multiple runs to allow them to happen exactly the same way. Three, I say have you actually kept the conditions of the test static over the period of time you were running the test? And normally that failed. In the third case especially because you often get the case where people say I ran the test but it didn't work very well. So I changed it. And now I got better results.

I said I started my life in physics. And here's a common thing every first year physics student does and I have seen it. You put a first year physics student in the lab and say here's your lab equipment, I want you to measure the speed of light. Here are your lasers. I want you to measure the speed of light. And the students work really hard and they have their lab notebook and write it on down. And they get the speed of light to 8 digits. And they say I got the speed of light. And they say how did you get that accurate? I tweaked it a little and adjusted the mirror until I got the speed of light. I knew the answer and I changed the experiment until I got it. Every single first year physics student has that experience. Unfortunately, a lot of senior researchers and information security still have that experience and change the conditions of the test in order to get the answer they are actually looking for. And I even see it in really rigorous environments. Where they really want to do the right thing. But for whatever reason the test equipment didn't quite work right and they weren't getting the results they expect. It's really similar to exactly the same thing you see.

Capturing and analyzing the data. After you've run the experiment and collected the results of the experiment. It's still not time to write the paper, right? But in many cases, this is where we see in the community is that I've done the experiment and collected the data. It's time to write the paper. When I get the conclusion section of the paper, I interpret what I think the data means. I'll interpret it at that point and understand what is going on there. And to identify problems as we look at it and if we have done the experiment right maybe I'll identify problems there and go back and rerun the experiment. By capturing and analyzing the data is really important forgetting this stuff out there. How many have been following the controversy on faster than light? Capturing and interpreting the data and giving it to people to identify is a really important step. You can come out with your claims and show it supports the data. But having it especially surprising claims be verified externally by other people and encourage people to do it. One of the nice things about that whole exercise which I've seen more in the security world was when the original scientist discovered the result and put it out there, they wrote we would like other people to verify this. This is something we don't quite believe. But we can't find anything wrong with our open rat is or technique or evident medication of error. We can't find where this is wrong. We have to tell you this is what we found. We will give anybody the data. All the data we collected. Please run it through your own verification. So that's exactly what happened. There were a number of places that first started with the data that was collected and try to confirm the same results and we had both. Both confirmed and not confirmed based on the same data set which I always find interesting. But in both cases people are trying to confirm the results. And we have other experiments that attempt to recreate the conditions.

I think this is the kind of process that I wish every single information security student had the experiment to be in that kind of roll or be taught that kind of a thing. Here we have a multi level secure system that should not allow information flow between these two levels. We have a result that looks like there is a potential band width between these two levels at this particular rate. Verify it, confirm it. Performance test that either confirms or refutes the claim that was actually made in that area. Again, one very small example and you say that's not what all information security is about. You are right. But doing that in a couple cases would help train students in the right direction. And help train where they need to go.

Reaching conclusions. This turns out to be a huge issue in observing papers and information security over the last 3 or 4 years. In many of the papers that I have personally evaluated, what I will see at the end of the paper, I'll get to a paper that looks if they attempted to apply a rigorous process. And I'll be saying this looks good. They are telling me how they put this together and how they ran the experiments. But they are doing really well. And they get all the way down and say at the very end and this cures cancer. This makes the internet secure. From this, we can determine everything we want to know about analysis and e-mail. People will reach conclusions way out of scope of what they had actually tried to confirm in their own hypothesis. So what we've done wrong is we clearly haven't helped create a community that knows how to appropriately draw detailed conclusions from the work and suggestions on what that work might lead to. Things that might indication. They can't show but could be perhaps the next step. This has really discouraged people from being incremental. What would encourage incremental advancement? I've come out with this result. This is where I can show given the tests that I have, it seems to indicate that this would work in a larger or slightly different context. But that was not in the scope for this particular paper. For one thing, they probably wouldn't get that paper published. The reviewers would say why didn't you do the other work too and try to get it out further? But if a paper like that were published, we could take it to the next step and see if it were confirmed.

Sharing the knowledge. Now we're writing the paper. We've gone through all this period of time to go through this and finally gotten to the point where we can write the paper and get the paper out there. So what's the problem here? First of all, it's hard to write a rigorous paper on a small incremental set. There are very few places that will accept a paper like that. Doesn't really meet the criteria. Not flashy enough. Doesn't think enough people would be interested. We don't really have a place to fit it. We don't have a good feel that allows us to say this fills the particular gap inside our knowledge. We don't have a good way of breaking the knowledge down. Over the past 3 to 5 years we decided threat is the best way to breakdown security which has led us in only looking at the small defensive stuff and not taking a more wholistic view. We've decided threat is a fundamental thing we break it down and we put everything down into how it addresses threat which is entirely responsive element to what we're doing. So applying the knowledge gained, putting it out there in a way everybody can sort of get to it. Is it a problem we haven't done that yet? No. We're a pretty young field. But it's time. It's time for us to say this is the fundamental thing that brings it all together. And I can create the periodic table and put things together. We need that level of insight. If I had that level of insight, this talk would be about that, not about this. I don't know what that structure is. I do know that we should have some of the smartest people in information security working on that structure, understanding how it works. I do believe that would be something we are explicitly funding and trying to make happen.

This actually led to the second. So that was the end of the first issue time. That's why things take too long all the way through publishing. The other two I'm going to spend less time on. The next one is peer review. So I'm sure most people in this room have been on conference committees, have been in program committees. Been in these various places where you are responsible for reviewing papers. Proposals. We have panel ss. All these kinds of things that are happening here. I have been organizing a number of these kinds of things as well. One of the things I find is it is really hard to find a good reviewer. And the handful of good reviewers there are are way overburdened reviewing things all the time. They are identified as one of the good reviewers that are out there. I will say most people who peer review a paper do it through a lens of would I do this work? And am I impressed with the results? And do I like the way this person wrote it up? Really comes down to that. And if I could, I would talk all those reviews out but that means I would talk 90% of my reviews if I did that. Getting good reviewers is hard. Why? Why is it so hard to find good peer reviews? Because we don't train people in information security to critically review work. We just don't do it. We spend all of our time teaching them how to program design systems, integrate, move forward. We don't spend time talking about critical review. We don't appreciate critical review when it happened. We don't consider it constructive. As a field. Now, I'm painting with a broad brush. Every one of you could say wait a minute I know a good reviewer and I know these people who can't wait to get a good critical review. True. I have some of those folks as well. It is true there are more of them today than there were three years ago. So we're moving in the right direction. But we have to forcibly continue to move in that direction. Right? We can't just assume we're always going to get more and more good reviewers. We have to attempt to do something to get more and more reviewers. In this kind of a world. So we love point solutions. We don't have the training. We cut parts of the page count. We find little value in incremental progress. All of these are internalized inside of peer reviewers. That's kind of the really tough thing and where do these peers come from? How do we gather our peer reviewers up? Social networking. Successful publications. Similar backgrounds. So this is a self perpetuating kind of culture. It's really hard to change. What makes this one of the really hard problems? The people who select the people who review papers are poor reviewers. And then we get a couple good reviewers and they only know a few good reviewers and those places turn out to get published in a lot of times they are not successful. New conferences, workshops, some of these kinds of areas. So we're organizing something called the national simposium on target research. We are attempting to find really good reviewers. So criteria was people who understand critical review. If you look at the program committee, you'll recognize names on there that are really good at critical review. I'm going to read this portion. Is there scientific evidence to show moving target techniques are a substantial improvement in the advancement of cyber systems? Scientific evidence is a key portion of the evaluation. In a small way we're trying to say I'm setting a bar. I'm creating reviewers really rigorous very critical. And I know as a result of this I'm going to have a very large tremendous General -- TREJETREJECTORY. This is where we want to go if we want to build the field and use papers you can build upon after publication. Why did I decide this is the right kind of thing to support? I've seen a lot of workshops with really good papers. I look at the papers and see if this paper had been peer reviewed and gone into a simposium, we actually build on it. And show this is part of the foundational research in the area. Why don't we try to put that body of knowledge together? Is moving target the right point area in the work to put it. Maybe, maybe not. But it's the best I've got in trying to find a niche that fits our agenda and makes scientific progress. The second problem again, peer review. So what problem does that have as a derivative problem? This is expectation of a breakthrough in every paper. So you go to the conference committee, you get all your papers in and start evaluating the papers. And what criteria do you use. A lot of times you'll say this was really supporting a breakthrough. It could be a really lousy paper and not support the breakthrough at all. But it claims to change the world. It's a really important paper. The higher the claim, the more important the paper. I think we've all seen areas in which we get caught up in this kind of a thing. The more important the research. This research is really important because it's widely applied. Again, that's not saying anything about whether the research was substantiated. That's saying was it a shiny object? Can we change that big shiny object? This is the next hot thing that's going to be out there. So a lot of times what this means is we equate scientific advancement within ovation. There's absolutely nothing wrong within ovation. My cell phone wouldn't be what it is without innovation. But to equate scientific papers within owe VAT I have papers is a mistake in terms of trying to put those two things together.

We also really talk about change rather than adding to our understanding. We make that as a higher delta. Things that change the change we think about stuff is better than stuff that adds to our knowledge. We somehow make that better. By the way, things that change, we don't put it through rigorous evaluation. We just think it's better. And of course, I've got enough gray hair to realize the field has a very short memory. If I see another paper that invents a detection, I think I'm just going to do something drastic.

We have very short memories in the field. I make students read papers from the 70s and they hate it but they are PDF photo copies of diagrams. That's the way we did it in the 70s. But that's where a lot of the fundamental stuff is. I'll have the students critique those papers and say how does what's in that paper relate to the problem today. And surprisingly, it's dead on with the problems we have today. Dead on.

So that's part of the problem we have there. So we don't separate this idea of engineering, innovation and science. And there's nothing wrong with any of these. Good engineering is good engineering. These are really good things. I would not want to have innovative bridges. I want good solid engineering. I want my word processor to have good solid engineering behind it. Cause I want it to work reliably all the time. But I want innovation because I do like shiny objects. I just don't think it's science. Shiny objects are cool. Take a risk. Maybe it will work. And sometimes innovation will lead to an observation that will allow us to make a causal relationship we never knew before. May happen more often than not. I look at social networking and innovation in social networking has led us to scientific questions about how people interact and how things are actually learned and how we actually move forward. Innovation can lead to scientific questions but it's not the same thing. Science is about why. Understanding the world we live in. Understanding better what's going on. Fundamentally knowing details in what increases security. What does the word increase even mean? We can't even get down to the idea of knowing what the scale is because we can't tease that out from innovation and engineering. How do we expand what we know. How do we create the lexicon? That's why it's important to overcome. This should not be about shiny objects and this should be. And we need both. But we need scientific advancement that doesn't focus on shiny objects.

Okay. I'm getting close to the end so I've got to quickly get to what are we going to do about this. I'm not going to say why we care. Cause I hope we all care. What are we going to do? General high level things. These are not very specific and kind of a composite of things I've heard over the years when I've heard people talk. Fund graduate student insists on completed experiments. I myself have fallen into many of these kinds of things and supported them. And these are all good things. But they kind of won't give us immediate accidentable things that we can do. Sort of the difference between a goal and objective. So these are good goals but don't give you short-term objectives you can operate to. And what are we going to do this year. I got a few other ideas that I want to put out there.

Leverage to people that do have clue. How do we leverage the people that have clue? Fund textbooks. Fund ways to capture the knowledge of the people that know how to do this well. Create something like a new academy module on scientific experimentation. Find another way to do it. Use other educational innovation to get the word out there. Recognize the few people out there that are doing a good job in this. Support them in terms of doing this. Education. This kind of goes into the more classic education. CS students are not trained in scientific method. Aside from textbooks, let's take a step back and say what are we training computer science to do? Maybe there should be a completely separate track for people who want to do engineering, innovation and science. Informally, there are three separate. Go to any major university and you'll see there are informal tracks based on whether they are interested in starting a company, whether you are interested in becoming a faculty member or doing science. Or whether you are interested in sort of becoming more rigorous on the information that we already know. If you really like building stuff based on what we already know, take computer engineering. If you want to do science, learn how to critique. Do some of the other things that allow you to become a scientist.

Simple thing. Put a $5,000 prize on the table at 2 or 3 major computer conference for the year. Do that 3, 4, 5 years in a row. I guarantee that will motivate people to write scientific papers. It's not enough to put the money on the table. You also have to put in this is the criteria for what a paper has to have. If it's an experimental paper, has to have a methodology. Theoretical paper, it better have a proof. And you better be able to stand behind the things that you say in that paper from your hypothesis down to your conclusion. You better show application of a rigorous scientific process if you want to be eligible for this prize. But that prize over a series of years which change the culture because students will write those papers. Students will actually begin to perform the science in terms of getting in place.

We're going to do your banquet. We'll do your annual banquet if in your call for papers you add scientific rigorous as part of your evaluation. This is not expensive. This isn't huge amounts of money. Considering how much money we put in student support and other kinds of things, we're not talking about coming up with a couple of new tens of billions of dollars to change the field. We're talking about relatively small and building out of that. Now, this is a good time for this right? Ten years ago we couldn't do that. Ten years ago prizes were not allowed. Couldn't do that kind of stuff. We have more and more evidence that that is successful in other areas within government. And it could be successful here.

Papers must use scientific rigor in their construction. How do we get there? Somewhere out there we have to show what that means. Which means somebody has to go through and go through the last 10 years and pull out a collection, and say these are the one that meet the criteria. Actually create a corporate that begins to point to papers that have already done this. I don't expect that to be big initially but I do expect it to exist. Structure the body of knowledge. Fund an effort now to actually try to discover what it is the different elements security. We need somebody who can actually support putting this kind of stuff together. I don't know who that's going to be. The person who knows how to do it probably doesn't know who it's going to be yet. But we have to have some real effort behind that to make it move.

Good data should be generated and cherished. This is a common theme we've made progress on. But we could do better. NIH requires if you are a publically funded work in genetic engineering, you publish your data for other researchers to use as a condition forgetting public money. Why couldn't we do that in computer security? Our data is sensitive. That's how we run and hide. Our data is too sensitive to that. I can think of sensitive stuff still required to go into the public data. There's lots of ways we can address this. We don't. Instead we back off and say we just won't do it then. We'll let everybody keep your own data and go back and ask. They won't give it to you. And we won't have public data like this. Gone a long way in trying to make progress but there's what we could do better. Tie the money that's producing the data that was demonstrated or consumed. That serves a lot of good. It allows you to reproduce results. It gives students good information to be able to actually do learning kinds of research to reproduce results. Three, it allows you to compare multiple different collection data collections against a similar hypothesis. It gives you a lot of good that can come out of this kind of work. We do need some focus there.

Needs to be explicit separation. Scientific and technological contribution. I'm not saying we shouldn't reward this. We do.

Making this a national priority. Now what do we do with that national priority? Some of these small items could help. Help to overcome those three barriers. I hypothesize that if we overcome those three barriers, we will have scientific improvement in the field. And that's kind of the bottom line.

Now, I can open up for questions.

And at this time if you would like to ask a question please pressstar 1. Make sure your phone is unmuted and record your name when prompted. Star 1 in one moment while we wait for that first question.

I have a question in the back here.

[Inaudible].

There have been classic collections of papers out there. Used to have a good question of this kind of stuff. I would characterize that as a pile of papers. Not really categorized very well. Not really shown to say these are ones you are going to read if you are going into this area. These are ones that maybe have been shown to not apply any more or have been overtaken by them. If we separate science to get into part of the question. If we separate science from innovation, I don't expect people to publish it on innovation. If, however, they are using public money to use science, I do expect that to be published. We don't insist on it today. I think we should. That's something we can actually do better and then add to the public body of knowledge. We're not adding to the body of knowledge in some key areas. People do see a proprietary.

[Inaudible].

Well, [Inaudible]. But yes, they do. I don't think they characterize the field. Yes?

So in your barrier, [Inaudible].

As far as we know.

I've seen proposals that say [Inaudible].

I don't think this is a barrier. I think it's an opportunity. I think there is an opportunity in the field to go beyond looking simply at a reactive defensive mode and take a view at understanding adversaries responses to our particular actions that we take in a much more inclusive kind of framework that would allow us to reason about what's going on, on the adversary side. But that's not a barrier to the science. That's an opportunity to understand the causal relationships. It does mean that it's harder. We don't have a determine is tick relationship and may look an awful lot more like psychology than it does physics. General, this is a thing we can embrace as part of the field, part of the understanding. Not as one of the three barriers. I think you can do science with a human adversary model. One of the things my day off is working is to understand how systems behave in the full context of what's going on, on both sides. Not just what's going on behind my firewall. But what's going on everywhere. That's part of what I'm trying to study is what's going on in the overall systems behavioral world.

Amazingly interesting question. As you say a lot of people shied away from treating in. So that provides opportunities like at the beginning of any kind of opened up area. So it's exciting much more than a barrier. To me that stimulates more work. That's why I didn't include in this as one of the bare juries to overcome -- barriers to overcome. Does that help?

[Inaudible].

I'm not going to try to summarize with the people online. I agree with all that. It's part of my bias. And I know I come from the world where causal relationships are absolutely well inclined. And that's kind of where we go. I would like to -- that's where my focus is going to be in approaching this field. I know it's not the only one and there are going to be other elements equally important. I will say this is for not something that because it's part of my bias, I believe that many of the three barriers I talked about here applied equally whether you pause at the sort of dependency that I'm worried about versus more of the economic style that are going to be out there that you might think of. I think the three live talked about still hold. But again, it's not my area as much in terms of looking at that. I would agree with you that my view is not the only view in this area in terms of defining. I'm not the right person to say here is exactly how the field should be structured. And I don't believe that because my strong bias really would push us down an area that I think is not inclusive. I think that's true.

[Inaudible].

Right. I completely agree.

[Inaudible].

My research that I've done over the past years has been based on the fact I can decompose the security problem into a few things that are manageable with regard to fewer variables. And that I can apply the result of that to some larger context. Now, in many cases I think that's been found to be true especially in the case protocol analysis and the more technical aspects of viruses and a number of things that worked well for. I have not focused my personal research on getting into what happens when we move into the intersection between the technology and the action of large groups. And I think that's where perhaps other techniques that I don't use myself are likely to be more affective. But that doesn't mean the approach I'm using isn't affective. Because I think there are individual elements we need to understand in the technical world that will lead us to questions and approaches even when we get to this larger world where that fundamental understanding helps. Understanding particle interactions doesn't directly lead to pressure volume temperature. But it does explain some kind of phenomenon by having both of these believes and holding them both true and coming into place. My suspicion is that's also going to be true in our field. Somebody could come along and say that's not true. I would read that paper with interest. Absolutely no question. I don't believe it at the moment. I believe the decomposition we are doing has real value and applying it to real context. You have to be careful not to draw conclusions beyond what the boundaries of your experimentation provides. I would never running an experiment showed something about virus propagation. Anything about Facebook propagation. That's completely different kind of a loadle. We're running 20 minutes late. I was told we can run 15 minutes late. I see we're losing people. I don't know if we have any questions from online.

We do have a question. Would you like to take that now?

Sure.

Sherry, your line is open.

HI. Thank you for this talk. It's great. I just have a couple comments. One is that sometimes when you talk to people about the need for a science of security they in many ways say I have a successful career and I'm getting publications so why should I change my behavior. We might try to think about that in terms of motivating people. And that leads to my second comment which is in security we're several decades behind what's going on with software engineering and there are some things that they've done that we might be able to embrace and make more particular to security and especially something like systematic reviews. In the medical community, the collaboration, that's evidence and does what's called systematic reviews. And Barbara has adapted that to software and engineering and a systematic review looks at the whole literature on a subject and say what do we know and what's the credibility of the evidence. So for example, [Inaudible] has done a systematic review of some studies on cost estimates and studies on requirements analysis. One of the things he's shown is there's a heavy bias if students evaluate that results. Now, we can do similar things in security and pick a narrow topic and say what do we know about virus propagation, what studies have there been, what have been the variables, where is the evidence confirming and where is it conflicting. That might address the first issue I raised which is to show somebody who is a researcher that there's value in having independent assessments. Or there's value in pursuing some variables rather than others or trying to clear up conflicting information from past studies. So that might act more as a motivator for people to start doing some of these more scientific studies.

Yeah. I think trying to find different motivators like that is really important. I fully understand the bias and some of the other things we can point out as well and I expect to see every bit of that as well. The way I usually think about this is how can we both as the community and as a funding box, a series of agencies and departments that fund research, what can we do that would begin to move this field in the right direction? Or maybe accelerate that movement? Which is kind of why I went to the recommendation. I think internally both educational institutions and other sorts of things might be motivators you are talking about with peer review that might be more appropriate for those people that have already embraced this is the right place to go. And if you ever have any ideas of how the two motivators could merge to multiply each other, I think that will be even better in terms of why it's going that way. At the end of the day probably both are required in terms of where we're going to go.

Do we have any other questions online?

At this time I'm showing no further questions.

Then in that case I'll thank you very much especially those of you that stayed to the end and hope to catch up with any of you who have future questions in the future. Thank you. [Applause].

And this concludes today's conference. Thank you for your participation. You may now disconnect. [Event Concluded]