Network Science and Engineering:

A Research Agenda

Jeannette M. Wing
Assistant Director
Computer and Information Science and Engineering Directorate
National Science Foundation

5 December 2007

In this short presentation, I will motivate one of computing's fundamental questions and I will challenge the community to articulate a research agenda to answer that question.

Let me start with some observations.

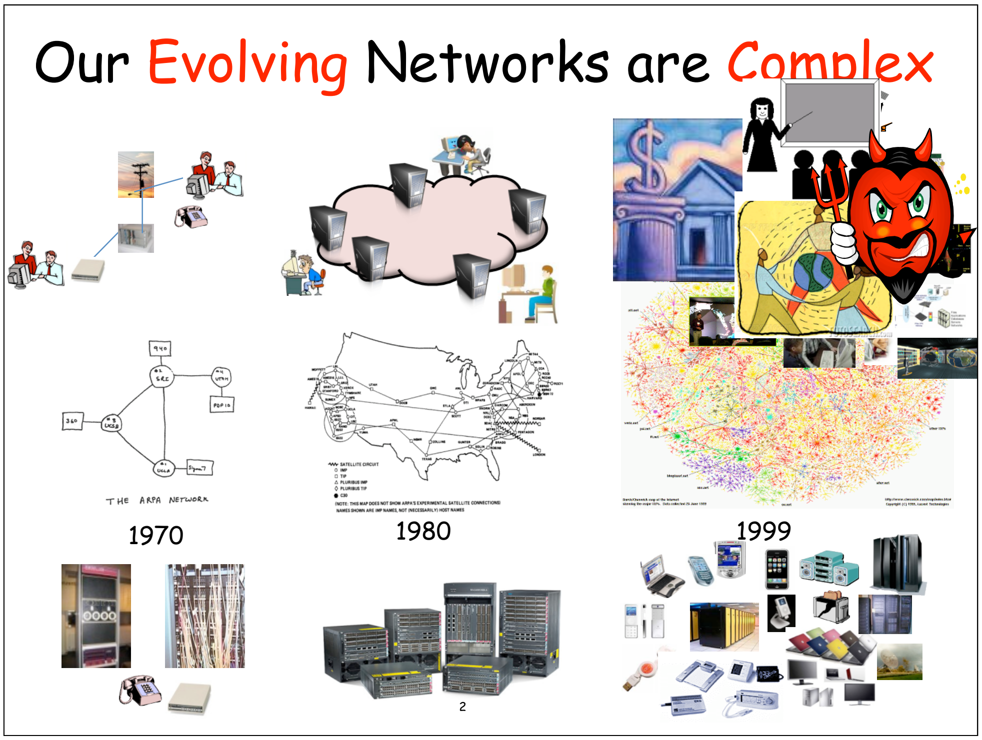

The networks that we have created and that have evolved over the decades are complex.

For example, the Internet is complex. Here is a picture of the ARPANET in 1970. Here it is in 1980. Twenty years later, it is too complex to draw, let alone understand, model, or predict its behavior. The Internet is computer science's gift to society, but ironically we cannot even describe it.

We interpret networks at multiple layers of abstraction. Below, we are concerned with new technologies: in the beginning we communicated via phone lines, modems, and cables underground and now we have a proliferation of communication media and a proliferation of devices, sensors, and actuators.

Above, we are concerned with new kinds of social uses of our computers and networks: from the days when people shared a terminal, to today where we have a multitude of applications, from the good to the bad. Examples of the good are on-line banking, social networks, and open courseware that already enables tens of millions of people, including tens of thousands of high school students, to learn from famous professors. Examples of the bad are spam, worms and viruses, and distributed denial of service.

I've shown pictures of the past and the present. What about the future? I believe that if we understand the complexity of our networks better then we can evolve them in ways that can unleash unimaginable creativity and innovation—from new technologies, to new applications, to new users—and hopefully at the same time improve the overall security of our networked systems.

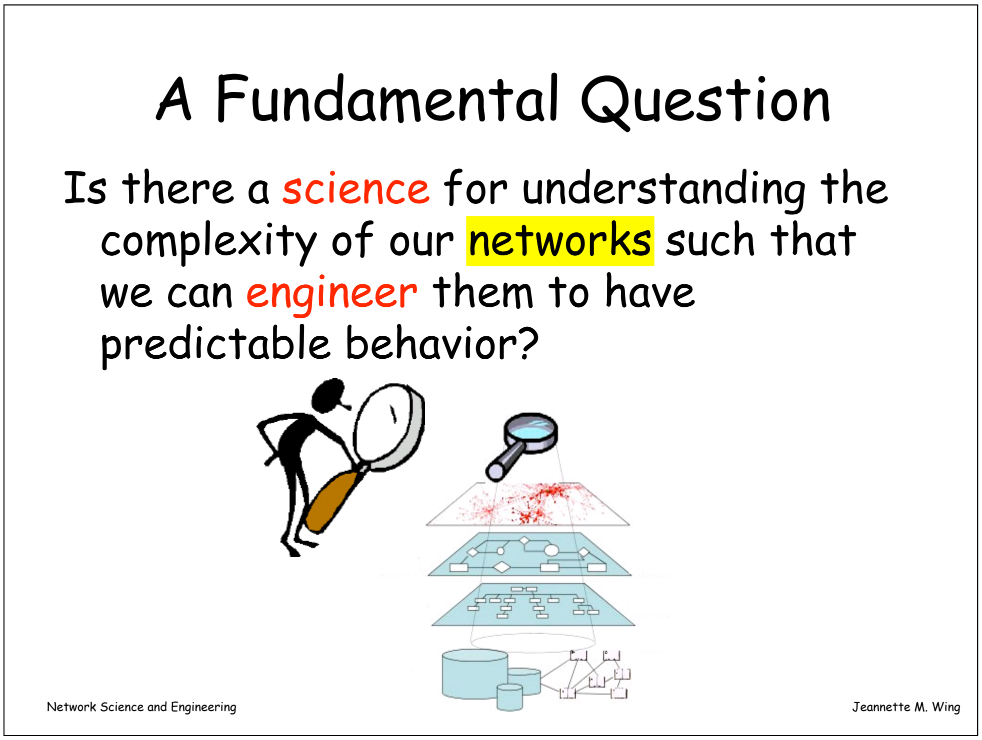

Motivated by these observations, herein lies a fundamental question:

Is there a science for understanding the complexity of our networks such that we can engineer them to have predictable behavior?

I interpret the term “networks” broadly. We need to understand networks at multiple layers of abstraction, as illustrated by the picture: from the physical layer on the bottom through multiple architectural and protocol layers in the middle all the way up to the top layer of networks of people, organizations, and societies.

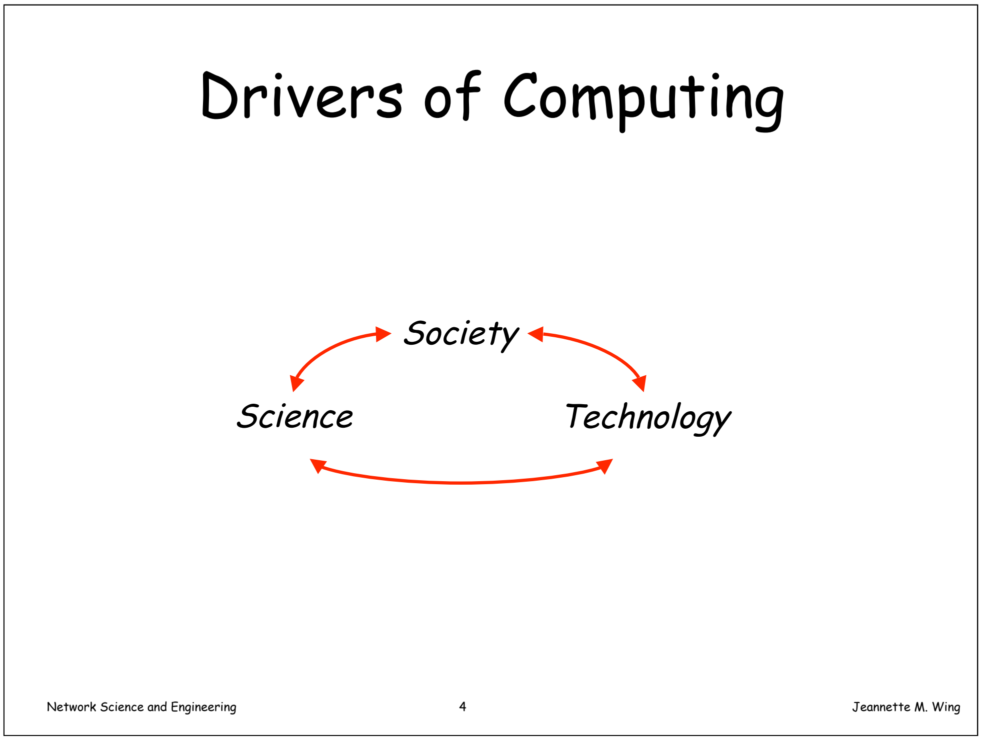

Interlude: I have been recently arguing that the field of computing is driven by scientific questions, technology innovation, and societal demands. I make this argument for two reasons: first, for ourselves in the field of computing, we are often so swept up by our technological advances or societal expectations that we forget that there are some deep scientific questions that underlie our field. Second, for others outside of computing, it is important to explain that the weight of each, and moreover the combination of our three drivers—science, technology, and society—make our field unique, indeed distinctive from other sciences, mathematics, and engineering. Why not celebrate this distinction?

Moreover, as shown by the bidirectional arrows, there is wonderful interplay—push and pull—among these three drivers: In the usual loop, scientific discovery feeds technology innovation which feeds new societal applications; in the reverse direction, new technology inspires new creative societal uses, which may demand new scientific discovery.

Underlying all three kinds of drivers to our field of computing, there are **fundamental, basic research** questions to be answered.

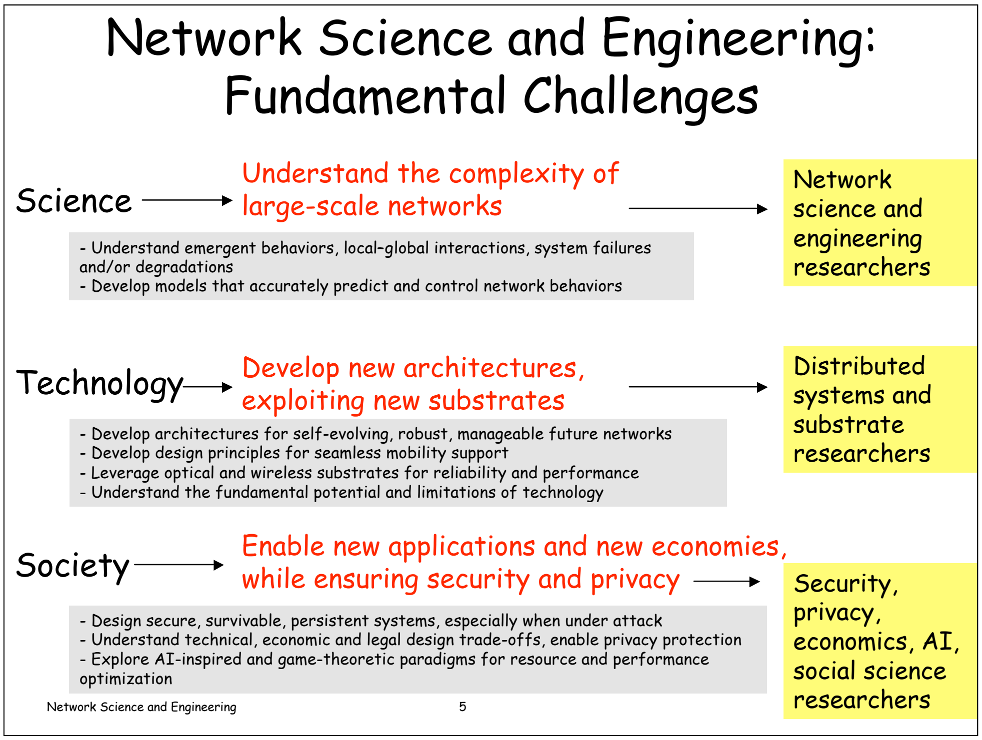

The field of computing has science drivers, technology drivers, and societal drivers. Network science and engineering is no different. For all three drivers, there are fundamental challenges that face us if we are to make progress in answering the question stated earlier. Let's use these three interacting sets of drivers to frame a research agenda for network science and engineering. I'll start with the science drivers.

The scientific challenge is to understand the complexity of networked systems. We do not have theories of our networks such that we understand their emergent properties. We do not have formal models of our networks such that we can assert any guarantees of reliability or survivability, especially in the presence of disruptive events or malicious attacks. We do not have models of our networks such that we can accurately predict their performance; for example Poisson models and heavy-tail distribution models are unrealistic or overly simplistic.

The technology drivers come from new communication substrates such as wireless and optical. These new substrates change the physical characteristics of the network, suggesting new network architectures, where the goal is not just simply optimizing the transmission of a packet from one host to another. For example, because of increased bandwidth and large distances between nodes we already view the network not just as a communication medium but as a storage medium. Technology drivers also come from new devices. The ubiquity of mobile computing and communication devices, for example, is forcing us to rethink our networks in truly fundamental ways, to include at the very least the human dimension of what networks are.

Moreover, new theories and new architectures enable new societal uses of the network—creating new opportunities such as virtual worlds and new challenges such as ensuring the security and privacy of users and their data.

For your interest, I've included in small font, a sampling of fundamental challenges for each of these three drivers.

The research communities active in addressing these intellectual challenges are correspondingly listed on the right. Indeed one of the opportunities in developing a research agenda for network science and engineering, is ambitiously to ignite a diverse range of communities to work together and look at networked systems from a holistic viewpoint—it's not enough to invent a new theory, a new architecture, or a new application without understanding the implications it has for all aspects of the networked system.

I would like to emphasize the expanse of these intellectual challenges. I am talking not just about networks as the invisible infrastructure that we all take for granted today, but about networks of people and organizations that creatively use the invisible infrastructure in unforeseen ways.

To calibrate the rapid growth in societal impact of the Internet, let me give you some concrete examples: Mosaic, the first browser, was created in 1993 and Google was founded in 1998, less than ten years ago. Google is known for search but it makes its money on ads; who would have guessed? (As as aside: remember that Mosaic and Google, like other browsers and search engines, were born out of NSF/CISE-funded research projects in universities.) Within just the past four years, we have seen the creation and influence of social networks embraced by the younger generation: SecondLife, FaceBook, and YouTube. We are also seeing old institutions remake themselves, including the print media like newspapers and the music recording industry. More sobering, in April of this year we witnessed the first cyber attack on a nation—Estonia—by a well-organized group in Russia. And just three weeks ago, the US company Seagate reported that hard drives sold in Taiwan contained Trojan Horses planted by subcontractors in China so that any information stored on those tainted hard drives was automatically uploaded to websites in Beijing. We did not anticipate any of these phenomena—good and bad—nor did we anticipate how they have completely changed the way in which people and organizations interact with each other.

It is the expanse of all these intellectual challenges that necessitates including not just researchers in computer and information science and engineering, but researchers in economics and the social sciences.

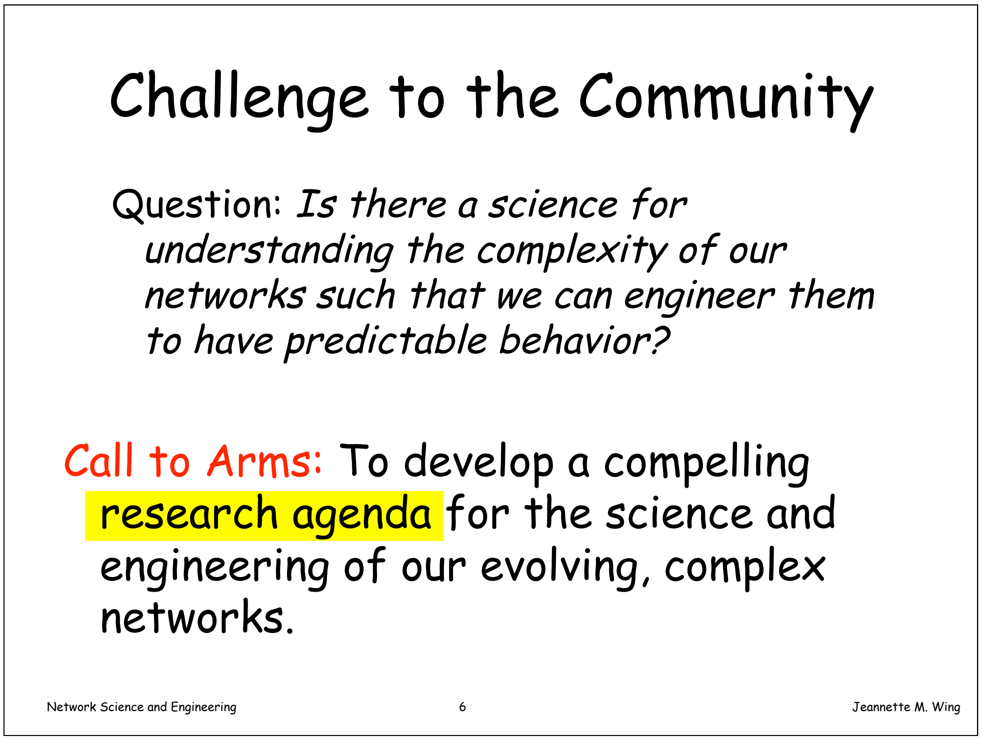

To recap, the fundamental question I pose to the community is:

> Is there a science for understanding the complexity of our networks such that we can engineer them to have predictable behavior?

Correspondingly then, the challenge to the community is to articulate a compelling research agenda for the science and engineering of our evolving, complex networks—at all layers of abstraction.

Before I close, I would like to make some personal remarks:

•I want to reemphasize interpreting "networks" broadly—to span multiple layers of abstractions, reflecting the role they are all anticipated to play in the future. Complexity arises not just from interactions within a single layer but also from interactions across layers.

•When I think about a science for understanding complexity, I mean it in the deepest and boldest sense: Is there a complexity theory for analyzing networks analogous to the complexity theory we have for analyzing algorithms? If we consider The Internet as a computer, what can be computed by such a machine? Let's call such computer a Network Machine, then much as we have a Universal Turing Machine, what is the equivalent of a Universal Network Machine?

•When I think about a way to engineer networked systems to predict behavior, I mean behavior in terms of both performance and correctness, and in fact, in terms of new dimensions of performance beyond time and space (e.g., power) and in terms of new meanings of correctness (e.g., notions to accommodate failures, dynamism, asynchrony, security and privacy, etc.); I mean behavior in terms of both quantifiable and qualitative measures. Most importantly, our understanding of behavior must reflect the dynamic, evolving nature of our networks. Maybe complexity and predictability are inherently contradictory?

To close, the vision for understanding the complexity of networked systems is audacious. Realizing this vision could revolutionize the future—as our networks affect every aspect of our daily lives, the efficiency, culture, economy, safety and security of our society. With NSF's support, we hope to give the community the courage to work in this challenging, new multidisciplinary context. The real work is ahead and rests with you: It is up to the community now to define a research agenda for realizing this grand vision.