Network Science and Engineering:

Call for a Research Agenda

Jeannette M. Wing

Assistant Director
Computer and Information Science and Engineering Directorate
National Science Foundation
and
President's Professor of Computer Science
Carnegie Mellon University

Stanford University, Palo Alto, CA
May 21, 2008

Our Evolving Networks are Complex

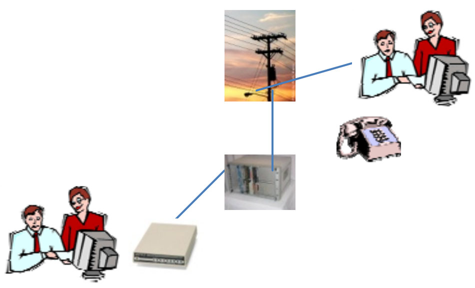

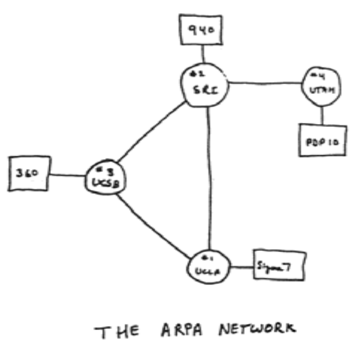

1970

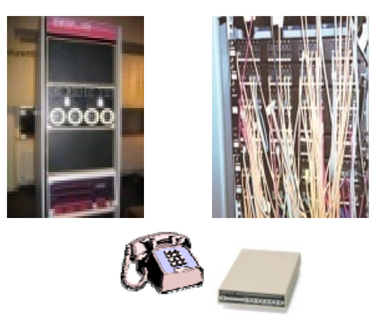

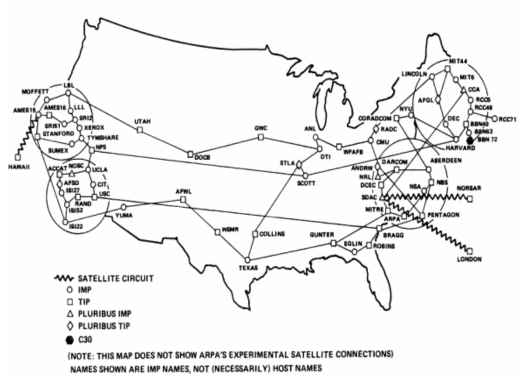

1980

1999

Challenge to the Community

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

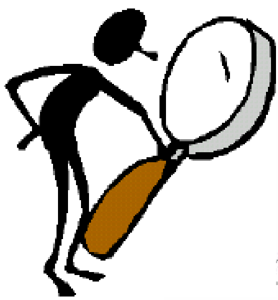

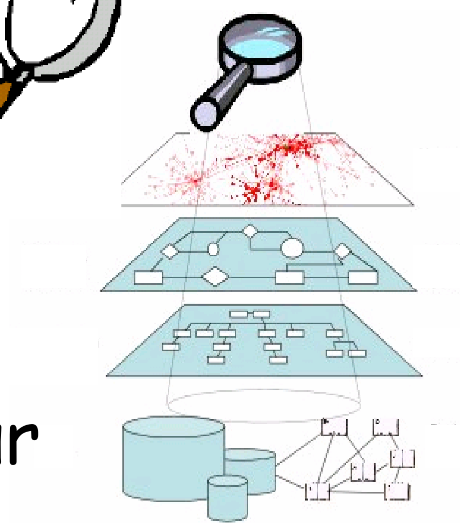

Call to Arms: To develop a compelling research agenda for the science and engineering of our evolving, complex networks.

Drivers of Computing

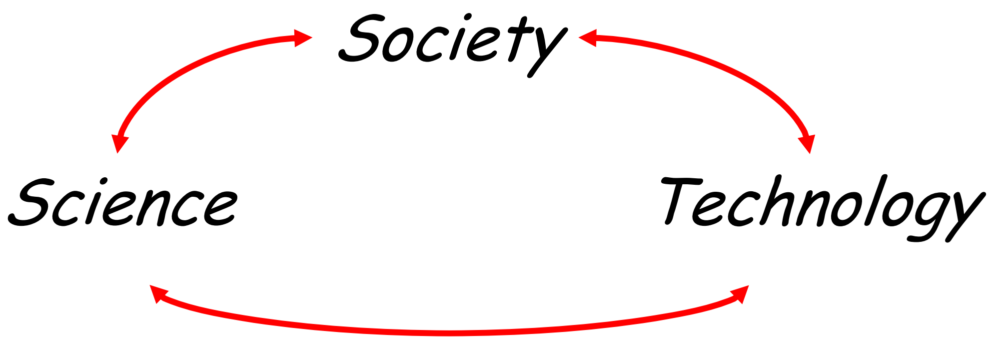

Network Science and Engineering: Fundamental Challenges

Science → Understand the complexity of large-scale networks → **Network science and engineering researchers**

- Understand emergent behaviors, local–global interactions, system failures and/or degradations
- Develop models that accurately predict and control network behaviors

Technology → Develop new architectures, exploiting new substrates → **Distributed systems and substrate researchers**

- Develop architectures for self-evolving, robust, manageable future networks
- Develop design principles for seamless mobility support
- Leverage optical and wireless substrates for reliability and performance
- Understand the fundamental potential and limitations of technology

Society → Enable new applications and new economies, while ensuring security and privacy → **Security, privacy, economics, AI, social science researchers**

- Design secure, survivable, persistent systems, especially when under attack
- Understand technical, economic and legal design trade-offs, enable privacy protection
- Explore AI-inspired and game-theoretic paradigms for resource and performance optimization

Complexity Cuts Across Abstraction Layers

- A societal pull may demand technological innovation or scientific discovery
  - Society ← Technology: tele-dancing
  - Society ← Science: energy-efficient devices, privacy logics
- A technology push can lead to unanticipated societal uses
  - WWW to Google to YouTube/MySpace/FaceBook
  - Small and cheap sensors, palm-sized devices, RFID tags

- Implication to the broad community
  - Working outside your comfort zone

A Fundamental Question

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

Characteristics of System Complexity

"Tipping Point"

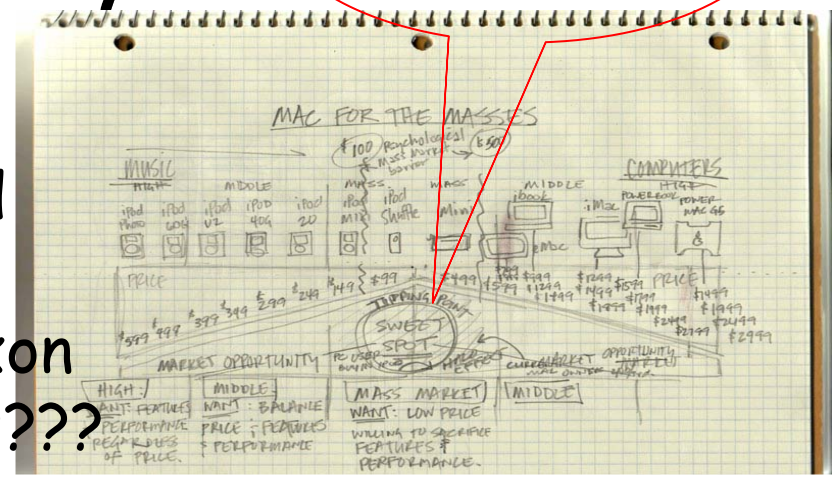

Tipping points

- Stampeding in a moving crowd
- Collapse of economic markets
- "Mac for the Masses" – P. Nixon
- 1970s: ARPAnet -> Internet ????

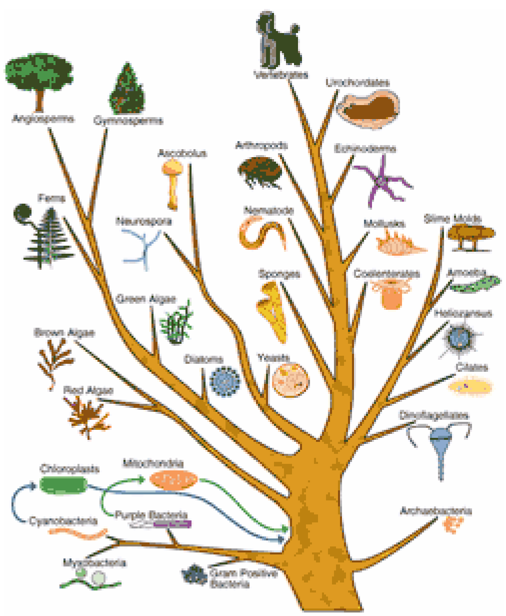

Emergent phenomena

- Evolution of new traits
- Development of cognition, e.g., language, vision, music
- "Aha" moments in cognition
- Spread of worms and viruses ????
- Open source phenomena ????

Predictable Behavior

- Predictable is ideal

A complicated system is a system with lots of parts and whose behavior as a whole can be entirely understood by reducing it to its parts.

A Car

A Car and Driver

A complex system is a system with lots of parts that when put together has emergent behavior.

Towards Predictable Behavior

- Behavior
  - Performance
    - Usual: time and space, e.g., bandwidth, latency, storage
    - New: power, …
  - Correctness
    - Usual: safety and liveness
    - New: resilience (to failure and attack), responsive
  - -ables
    - Adaptable, evolvable, measurable, …
  - Quantifiable and qualitative measures

- Most importantly, our understanding of behavior must reflect the dynamic, evolving nature of our networks

Sources of Network Complexity

- Inherent
  - People: unpredictable at best, malicious at worst
  - Mother Nature: unpredictable, unforgiving, and disrup
- Scale, in terms of
  - numbers of, sizes of, types of elements (e.g., users, nodes, connectors), and recursively, … of networks
  - distance and time, also at different scales
- Design
  - Mismatched interfaces, non-interoperability
  - Unanticipated uses and users
  - Violation of assumptions as environment or re
  - Lack of requirements

Food for Thought

- (Where) does the current Internet embed assumptions of plenty?

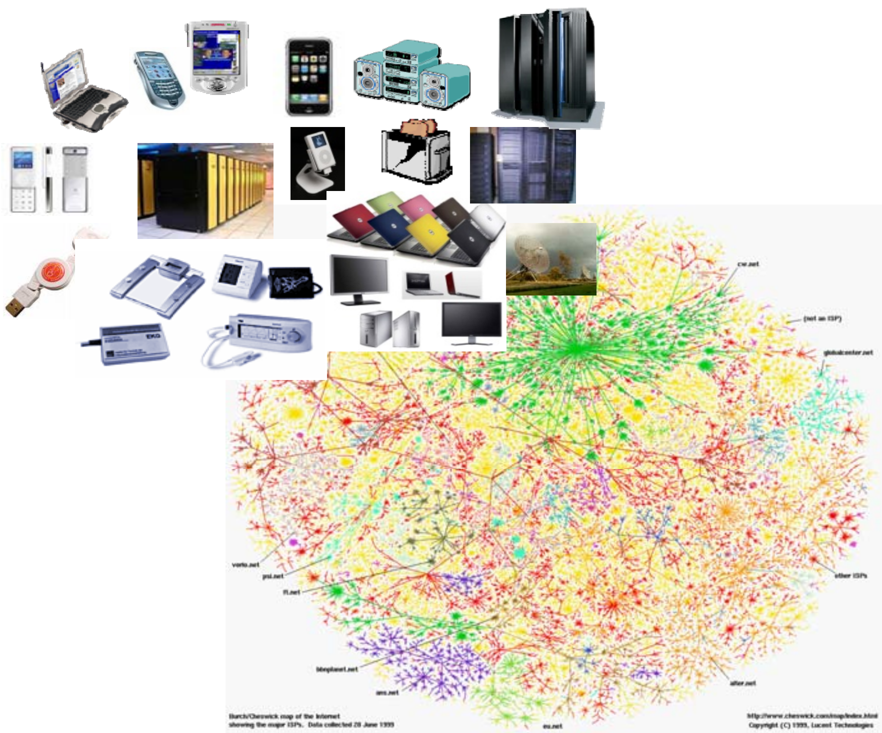

Does TCP work here? (Hint: no!)

Network Models

- Poisson, heavy-tail, self-similar, chaotic, fractal, butterfly effect, state machines, game theoretic, disease/viral, …
  - We know some are wrong or too crude
  - We are trying others
  - None consider all "usual" performance and/or correctness properties at once, let alone new ones
  - Composable models, e.g., per property, would be nice

- Maybe our networks are really different from anything anyone has ever seen (in nature) or built (by human) before
  - Implication: A BRAND NEW THEORY is needed!

Food for Thought
(courtesy John Wroclawski)

Electricity: 1800…

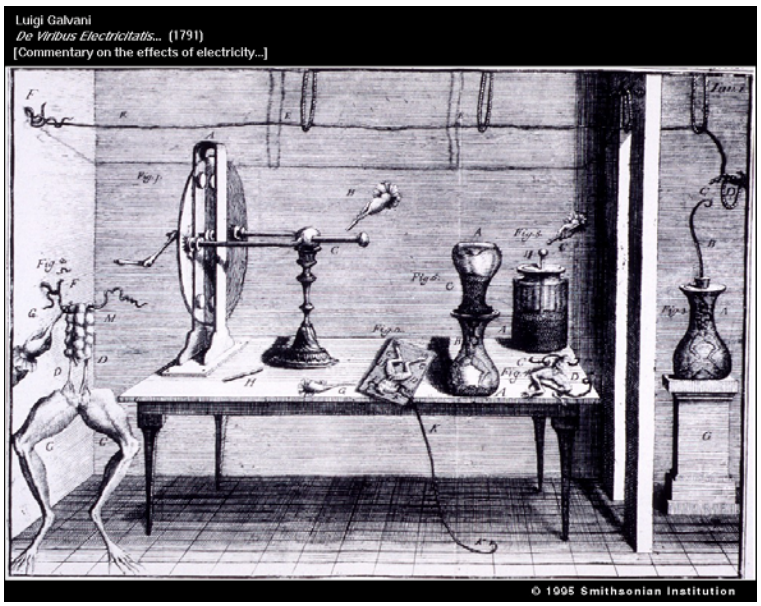

Electricity: Today…

$$\oint \vec{E} \cdot d\vec{A} = \frac{q}{\varepsilon_0}$$

$$\oint \vec{B} \cdot d\vec{A} = 0$$

$$\oint \vec{E} \cdot d\vec{s} = -\frac{d\Phi_B}{dt}$$

$$\oint \vec{B} \cdot d\vec{s} = \mu_0 i + \frac{1}{c^2}\frac{\partial}{\partial t}\int \vec{E} \cdot d\vec{A}$$

What are the analogies…
… for Network Architecture and Design?

Beyond Computer Networks

Utility networks
e.g., electric power

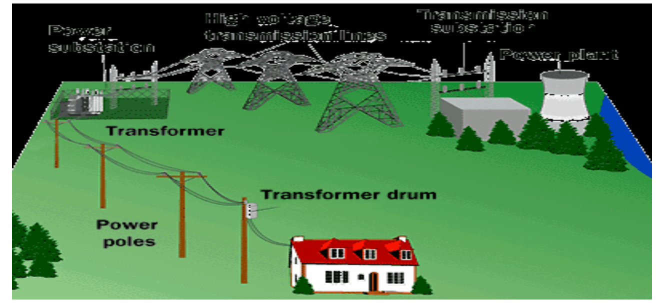

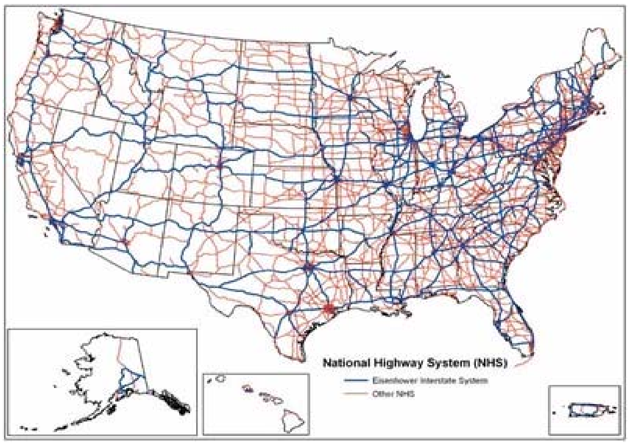

Transport networks
e.g., for cars, trains

Economic networks
e.g., a community of individuals affecting (
market

Social networks
e.g., friends, family,
colleagu

Political networks
e.g., voting systems

Understanding Complexity

- 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?
  - What is computable? [From J.M. Wing, "Five Deep Questions in Computing," CACM January 2008]
- 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?
  - Challenge to us: Could we build one?

What-if Applications

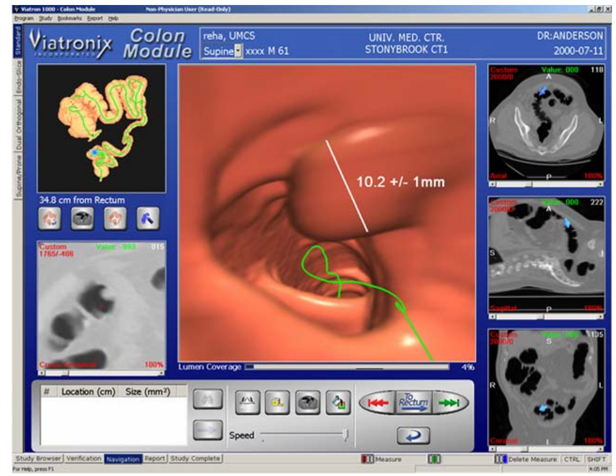

Five-sensory tele-presence, e.g.,
- tele-meetings (social aspects)
- tele-surgery (safety critical)

Ask anyone anything anytime anywhere

Automated vehicles on automated highways

Secure and private communication and data for all

Modeling the earth, modeling the brain

From Agenda to Experiments to Infrastructure

- Research agenda
  - Identifies fundamental questions to answer
    - aka the "science story"
  - Drives a set of experiments to conduct
    - to validate theories and models
- Experiments
  - Drives what infrastructure and facilities are needed
- Infrastructure could range from
  - Existing Internet, existing testbeds, federation of testbeds, something brand new (from small to large), federation of all of the above, to federation with international efforts

Feedback Loop

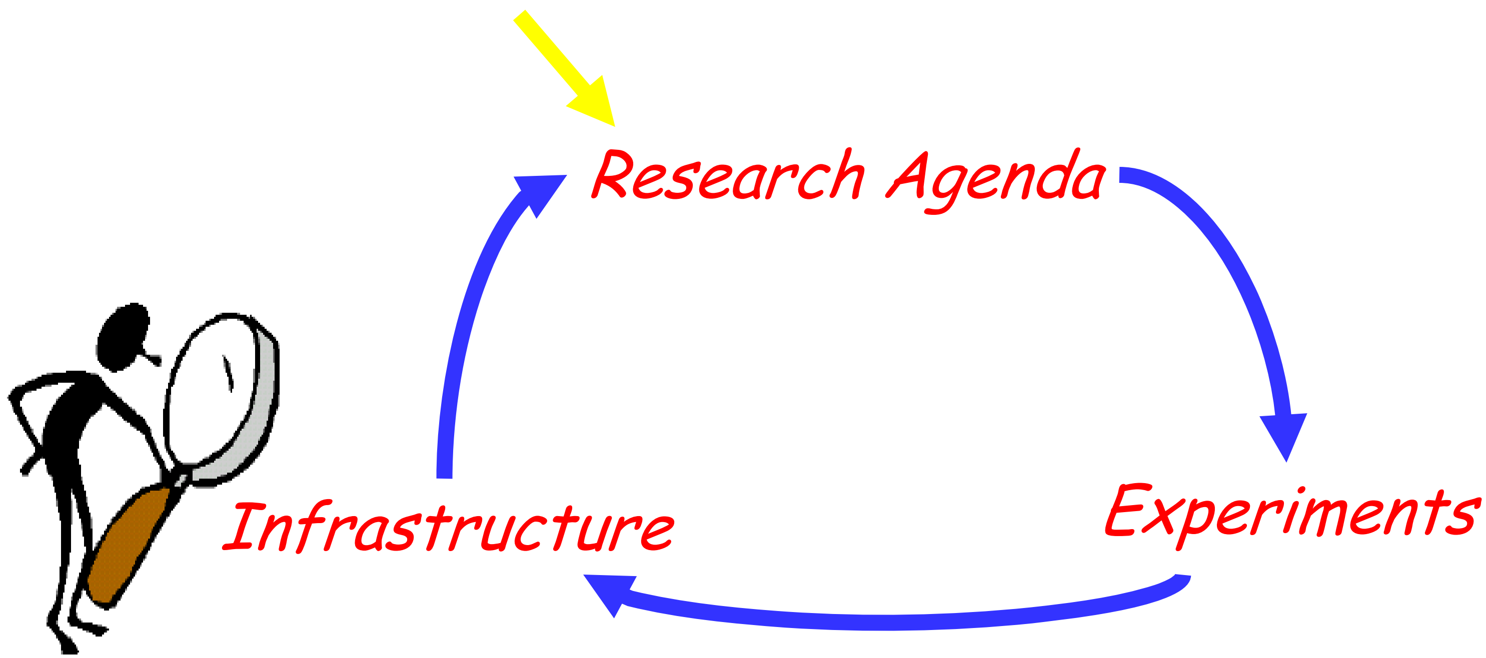

Prototyping the Infrastructure Needs

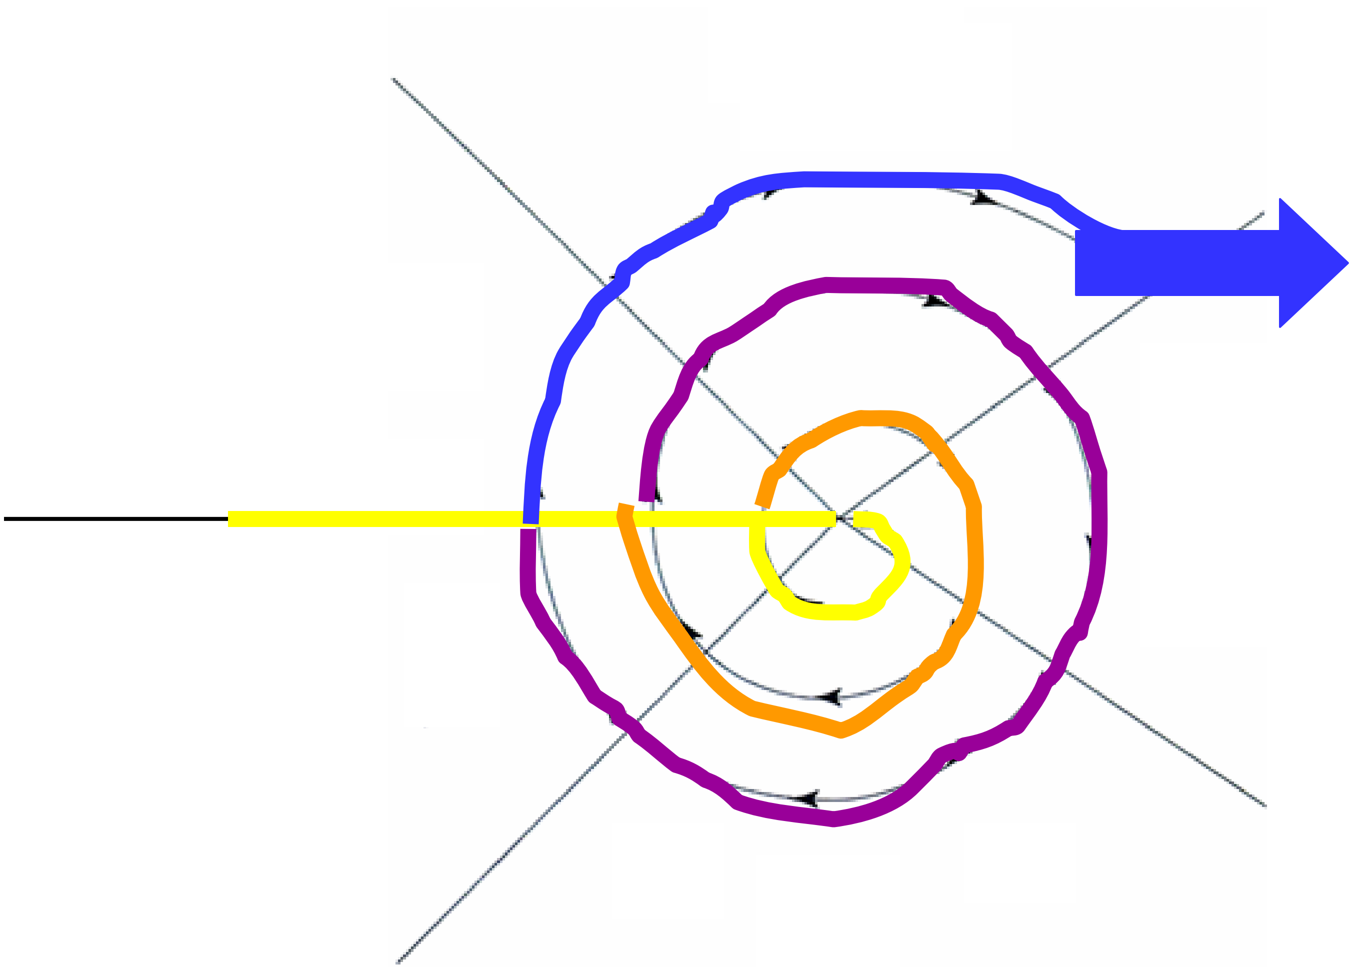

Suite of experimental infrastructure capabilities

Core Concepts

- Programmability
- Virtualization
- Slice-based experimentation
- Federation

Core Concept 1: Programmability (P)

Researchers can program the physical substrate, the means to transfer data, routing, how connections are managed, syntax and semantic rules, and how applications are serviced.

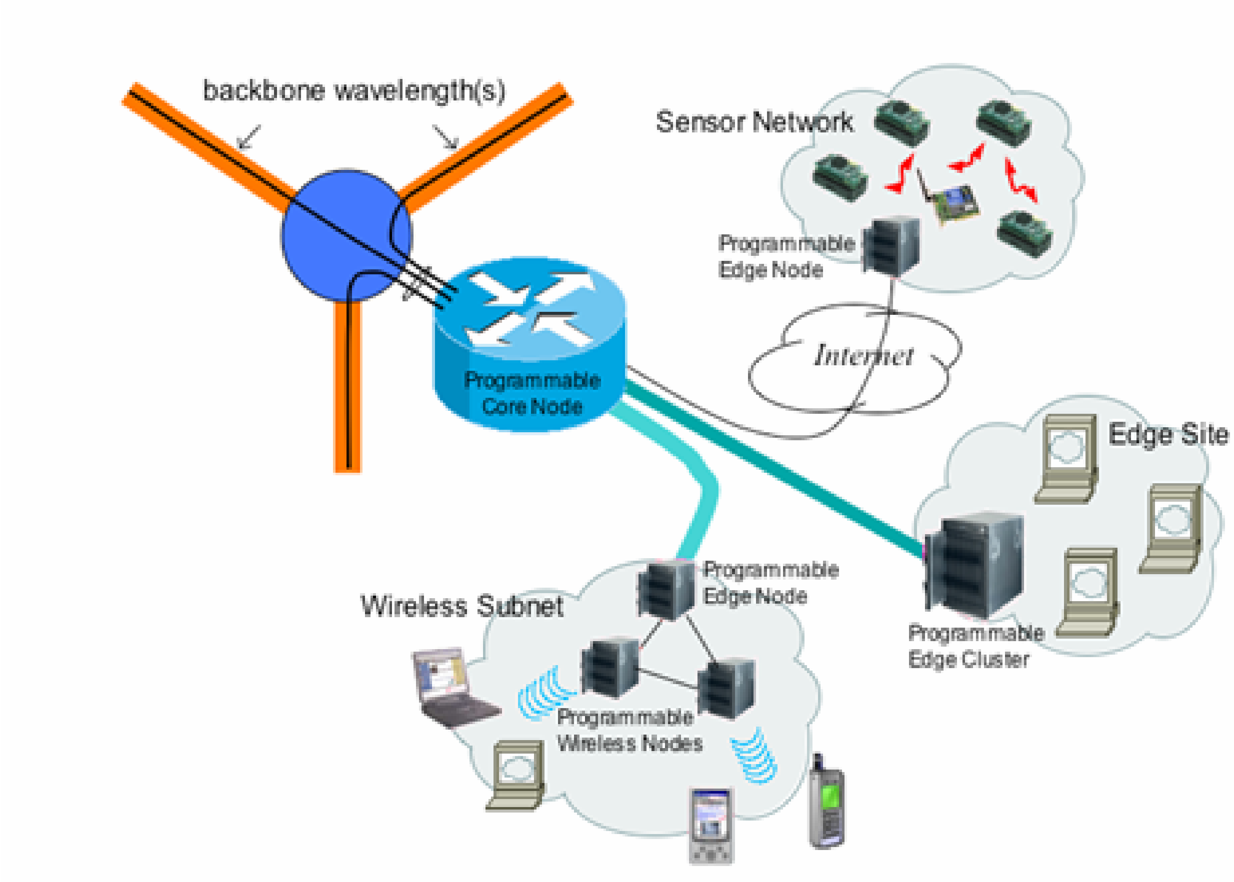

Core Concepts 2&3: Virtualization (V)

- Virtualization allows many different, simultaneous experiments.
- Each researcher has a "slice."

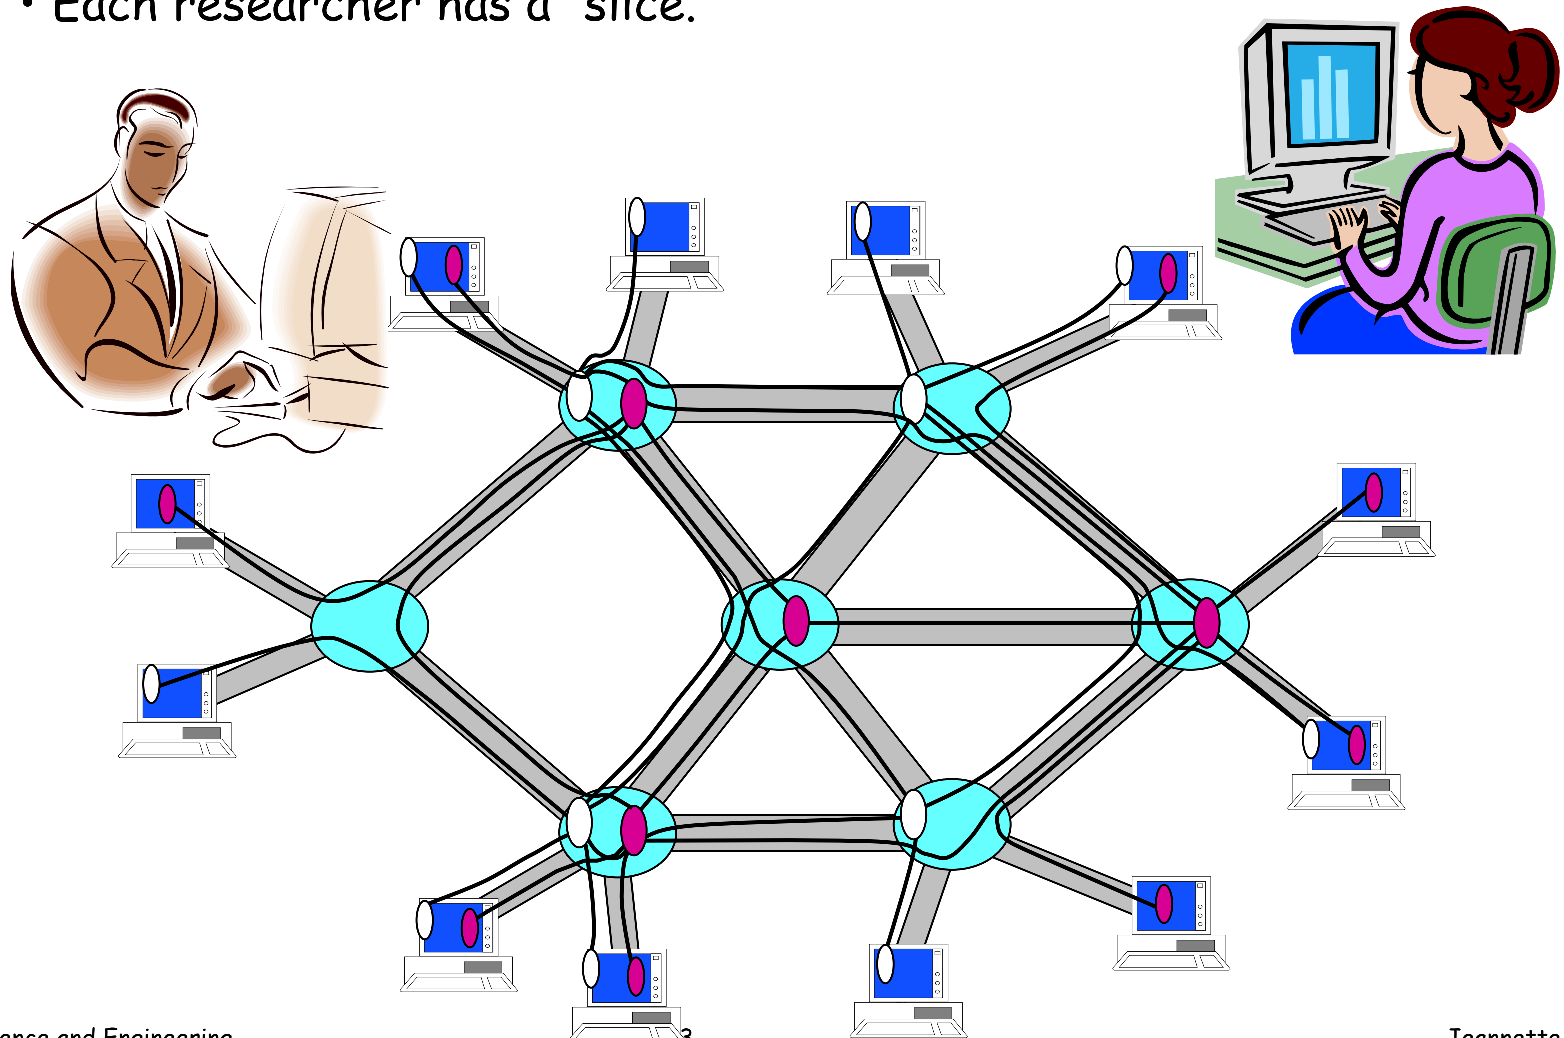

Core Concept 4: Federation (F)

- A federated infrastructure evolves by gluing together heterogeneous components over time, including emerging technologies.

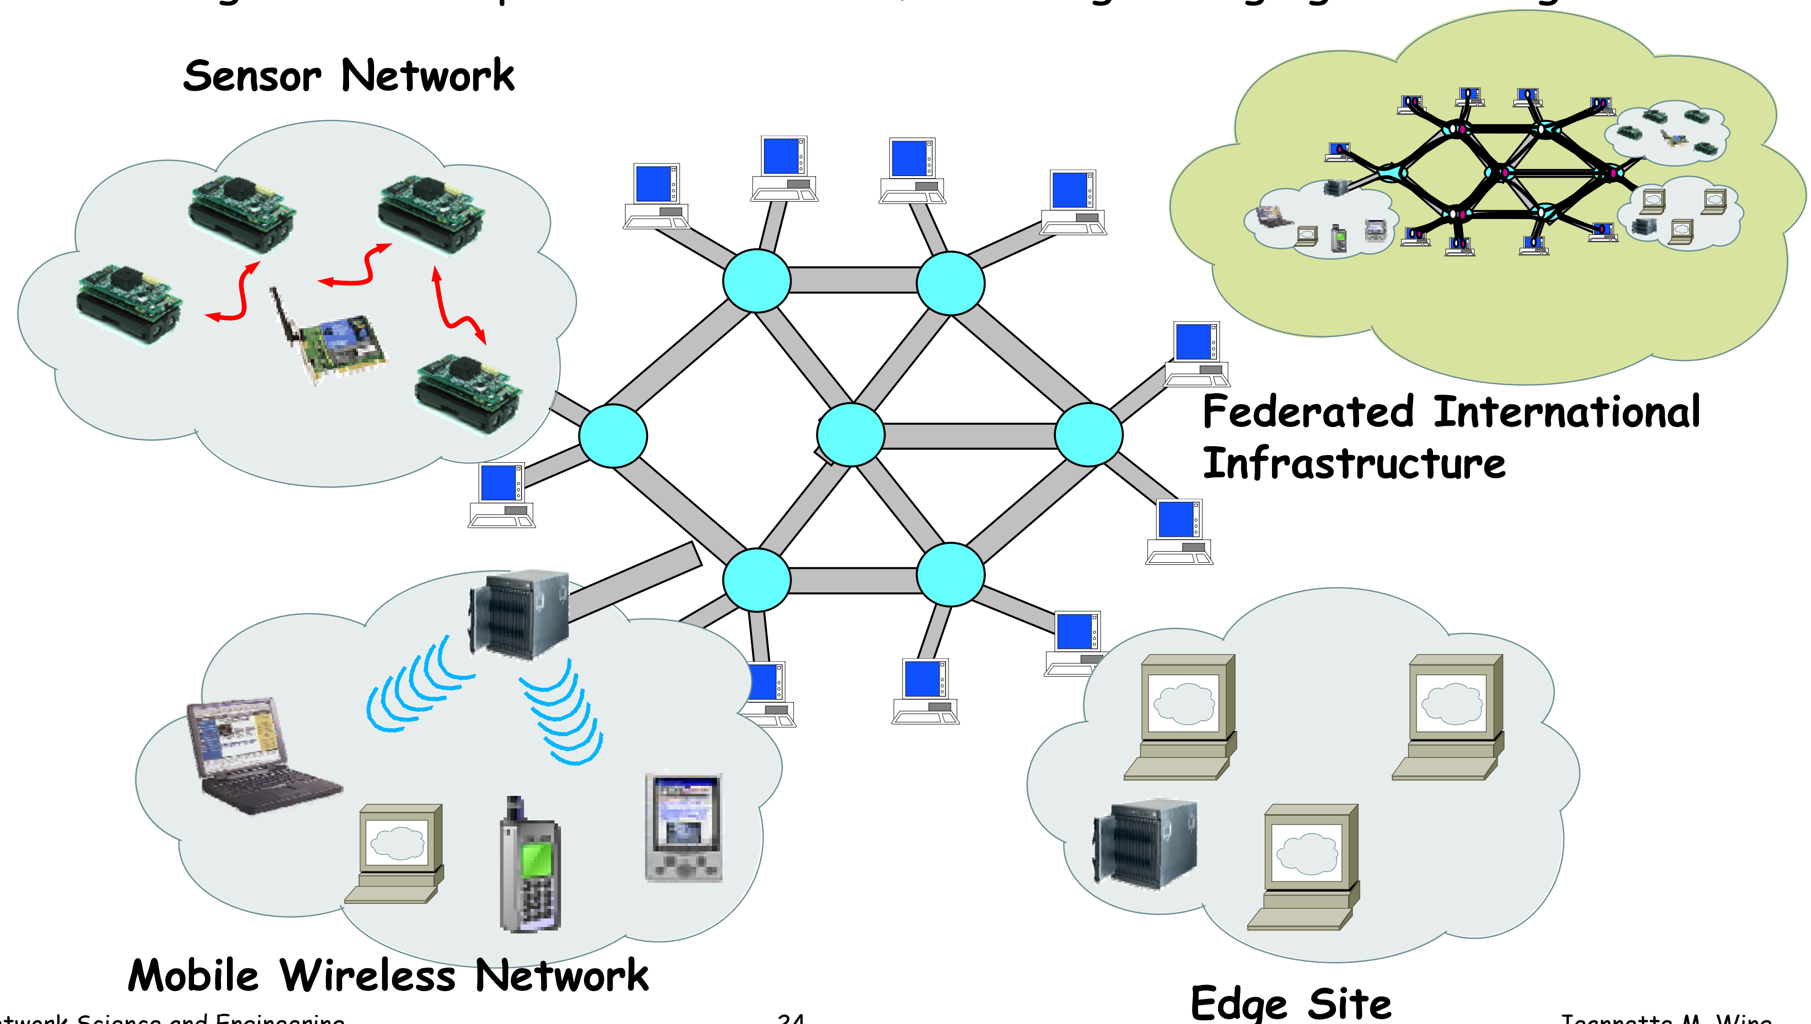

Secret Weapons

Exploiting Computing's Uniqueness

- Software is our *technical* advantage
  - Plus: We can do anything in software
  - Minus: We can do anything in software
- Unlike other sciences, prototypi ur *process* advantage
  - Feasibility – sa
  - Possibility – spo
- Implications iqueness
  - Power of softw s the nature of our infrastructure is different
  - Power of prototy g implies the nature of our infrastructure building process is different
- We are breaking new ground at the NSF!

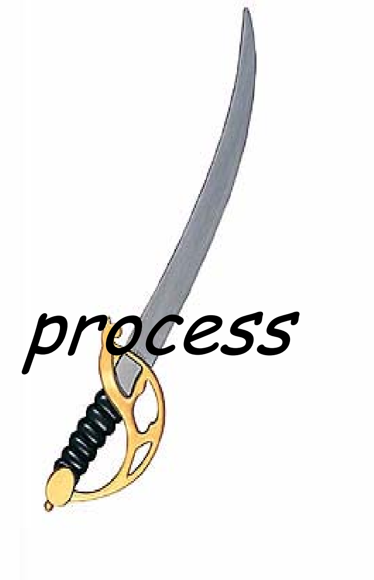

People

- NetSE Council: Ellen Zegura, chair
- GENI Proje… and BBN
- CCC: Ed La…, Sue Graham, …
- NSF team

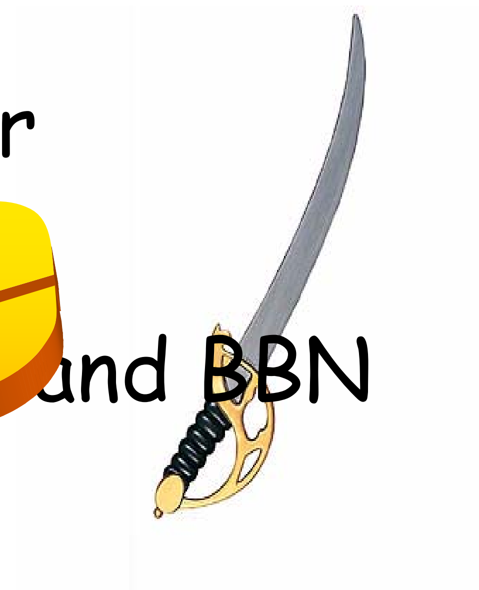

Synthesizing Discussion: NetSE Council

Mission (work in progress): The primary mission of the Network Science and Engineering (NetSE) Council is to articulate a compelling research agenda for Network Science and Engineering, including inter-related theoretical, experimental and societal aspects.

- Ellen Zegura, chair
- Tom Anderson, Washington
- Hari Balakrishnan, MIT
- Joe Berthold, Ciena
- Charlie Catlett, Argonne
- Mike Dahlin, UT Austin
- Chip Elliot – GPO (ex-officio)
- Joan Feigenbaum, Yale
- Stephanie Forrest, UNM
- Roscoe Giles, Boston Univ
- Jim Hendler, RPI
- Michael Kearns, UPenn
- Ed Lazowska, Washington
- Peter Lee, CMU
- Helen Nissenbaum, NYU
- Larry Peterson, Princeton
- Jennifer Rexford, Princeton
- Stefan Savage, UCSD
- Scott Shenker, ICSI/Berkeley
- Al Spector, Google

Existing Input

- Clark et al. planning document for Global Environment for Network Innovations
- Shenker et al. "I Dream of GENI" document
- Kearns and Forrest ISAT study
- Feigenbaum, Mitzenmacher, and others on Theory of Networked Computation
- Hendler and others in Web Science
- Ruzena Bajcsy, Fran Berman, and others on CS-plus-Social Sciences
- NSF/OECD Workshop "Social and Economic Factors Shaping the Future of the Internet"
- Current NSF "networking" programs
  - FIND, SING, NGNI

3 Focused Workshops

- Summer and Fall
- Focused at the “edges”

1. Science of Network Design
   - Co-chairs: John Doyle, CalTech;John Wroclawski, ISI
2. Network Economics
   - Co-chairs: Mike Kearns, UPenn; Colin Camerer, CalTech
3. Network Design and Societal Values
   - Co-chairs: Helen Nissenbaum, NYU; David Clark, MIT

CCC Charge to NetSE Council

1. Develop a compelling research agenda
2. Identify a set of experiments to carry out the agenda, e.g., to test out new theories/models of network complexity
3. Identify suite of experimental infrastructure needed to carry out the experiments
4. Work with the GENI Project Office in prototyping activity for experimental infrastructure.

GPO Activities

- GENI Project Office: Chip Elliot and team at BBN
  - Hard work in short period of time
    - Organizing and challenging the community to push the frontiers of experimental infrastructure
    - Engineering Conferences, Infrastructure Prototyping Competition (underway)
    - Working with industry and international partners
    - Establishment of working groups
- Working Groups: Architects and designers of the experimental infrastructure
- Community participation in working groups is welcome and encouraged!

Starting Exercise

- Identify existing testbeds and their features.

feature

testbed

- Identify experiments and their needs
- Do mapping
- What's missing?

Breaking New Ground *Together*

- Unexplored territory in network science and engineering
  - Broad scope for research agenda
  - New relationships among theoreticians, experimentalists, and systems and applications builders
  - New relationships with social science, law, economics, medicine, etc.
- Big Science is new for Computer Science
  - Science at scale, experimental settings at scale, real users at scale, user opt-in at scale
  - Scientists, engineers, technicians, managers, and funding agencies *must work together*

Challenge to the Community

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

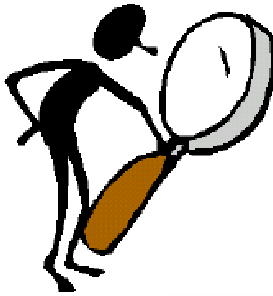

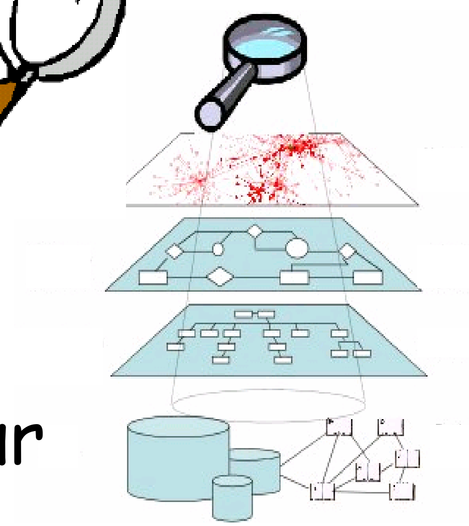

Call to Arms: To develop a compelling research agenda for the science and engineering of our evolving, complex networks.

Thank you!