# Goals for GENI

John Wroclawski

USC/ISI

- This talk lays out, at a high level, what we hope GENI will accomplish
- Throughout the day, you should ask yourself (and each other…)
  - Are these the right things to accomplish?
  - Does the facility you’re hearing about accomplish these things?

# GENI

- GENI is a facility to support open, large-scale, realistic experiments.
  - A physical infrastructure in which many different kinds of networks and distributed systems can easily be embedded.
  - A management framework (operating system) for this physical infrastructure
  - A collection of services, tools, and related material to make the infrastructure useful.

# One goal, two activities

Central goal of GENI:

- Change the nature of networked and distributed systems design by catalyzing and integrating
  - Rigorous theoretical understanding
  - Compelling and thorough experimental validation.

By supporting two activities

- Rich and realistic *controlled experiments*
- Functional and cost-effective *prototype systems*

# Intellectual Merit

- **[Science]** GENI will allow us to experimentally answer questions about complex systems, giving us an increased understanding about their dynamics, stability, evolvability, emergent behaviors, and related matters.
- **[Architecture]** GENI will allow us to evaluate alternative architectural structures, and reconcile the contradictory goals a network architecture must meet.
- **[Engineering]** GENI will help us evaluate engineering tradeoffs, and test theories about how different architectural elements might be designed.

# Broader Impact

- Lead to artifacts that provide value to society
  - An Internet for tomorrow
    - more secure, available, manageable, usable
    - suited for computing in the next decade
  - Catalyze the distributed digitized world
    - personal control of your personal life
    - real time sensing of the planet
- Lead to artifacts that provide value to science
  - Enhanced services that improve the scientific process

# Creating impact: the integrated research cycle

Goal: Seamless conception-to-deployment process

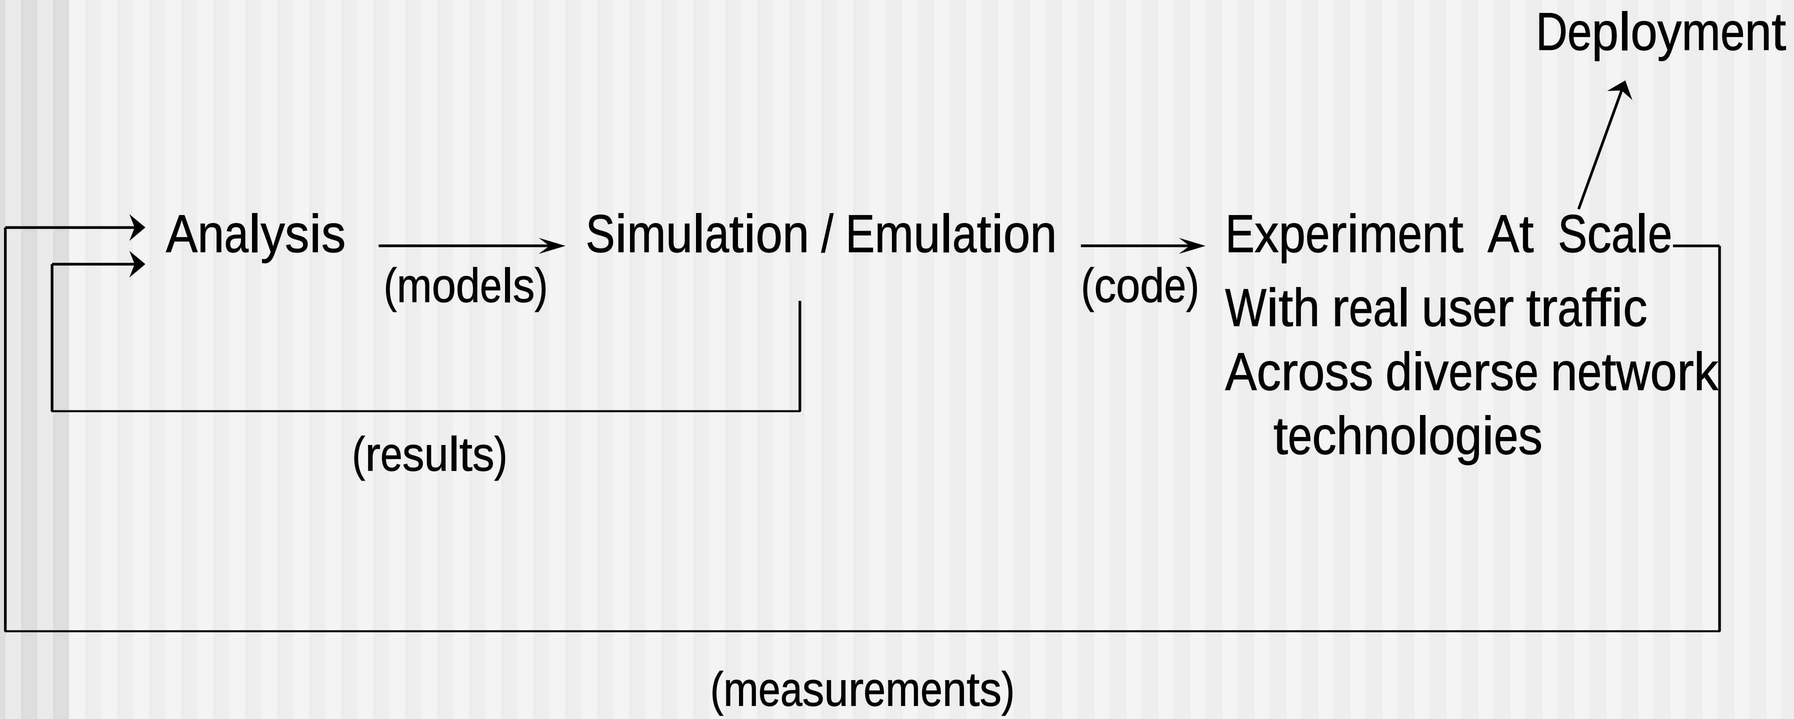

# Why this matters

- Systems-oriented CS research needs to build and try out its ideas to be effective. We're not smart enough yet to do anything else.
  - Paper designs are often idle speculation.
  - Simulation is only occasionally a substitute.
  - Misapplied experiments can mislead deeply
- We need:
  - Real implementation
  - Real experience
  - Real users
  - Real measurements
  - <u>To learn from the past and present, but live in the future</u>

# The past: some examples from networking

- DartNet (a DARPA program)
  - Validated IPv6, multicast, QoS, other stuff.
  - Hard to program and deploy, but a great success.
  - Other “similar” projects missed the mark.
- Gigabit Testbeds (a NSF program)
  - The thrill and peril of speed.
- More recent facilities
  - DETER, ORBIT, WAIL, …
- Overlays on the Internet
  - RON, X-Bone, PlanetLab, …

# Limitations..

- Simulation based on limited models
  - topologies, admin policies, workloads, failures…
- Emulation (and “in lab” tests) are similarly limited
  - only as good as the environment
- Traditional testbeds are targeted
  - often of limited reach
  - often with limited programmability (built in assumptions)
- Testbed dilemma
  - production: real users but incremental change
  - research: radical change but no real users

# The tao of GENI

- Significantly reduce these limitations…
- …over a broad range of research
  - The future of the Internet (FIND)
  - Technology innovation
  - Distributed systems and application support
  - New applications
- By enabling experiments where physical distribution, relevance of locality, scope, scale, long life, and access to real users matter.
- But *not* “any experiment”

# FIND: An NSF challenge question

- 1) What are the requirements for the global network of 10 or 15 years from now, and what should that network look like?
- To conceive the future, it helps to let go of the present:
- 2) How would we re-conceive tomorrow’s global network today, if we could design it from scratch?
  - This is not change for the sake of change, but a chance to free our minds.

# Q1: Isn't today's net good enough?

- Must start with serious discussion of requirements:
  - It's not just about cool new apps.
- Security and robustness.
  - As available as the phone system
  - Been trying for 15 years--try differently?
- Easier to manage.
  - Really hard intellectual problem
  - No framework in original design.
- Recognize the importance of non-technical considerations
  - Consider the economic landscape.
  - Consider the social context.

# Q2: What’s next?

- Service in times of crisis, and disaster recovery.
  - Resilience, heterogeneous distributed fallback, automated management of highly distributed systems, robust trustworthy dissemination of information
- Understanding the planet in real time.
  - Operation of highly distributed sensors.
  - Generation, control, ownership, privacy, use of global distributed information.
- Managing the global interactions of the personal experience.
  - Information, personal sensing, location, privacy
- Yes?

# Hence the scope..

- New science, new theory
- Impacts of new technology
  - Software defined radios
  - Programmable optics
  - …
- New computing paradigms
  - The Internet grew up in a stable “PC” time.
  - Tomorrow: embedded processors, sensors, everywhere
- New applications and distributed infrastructures

# A “New Theory” example: Theoretically Derived Architectures

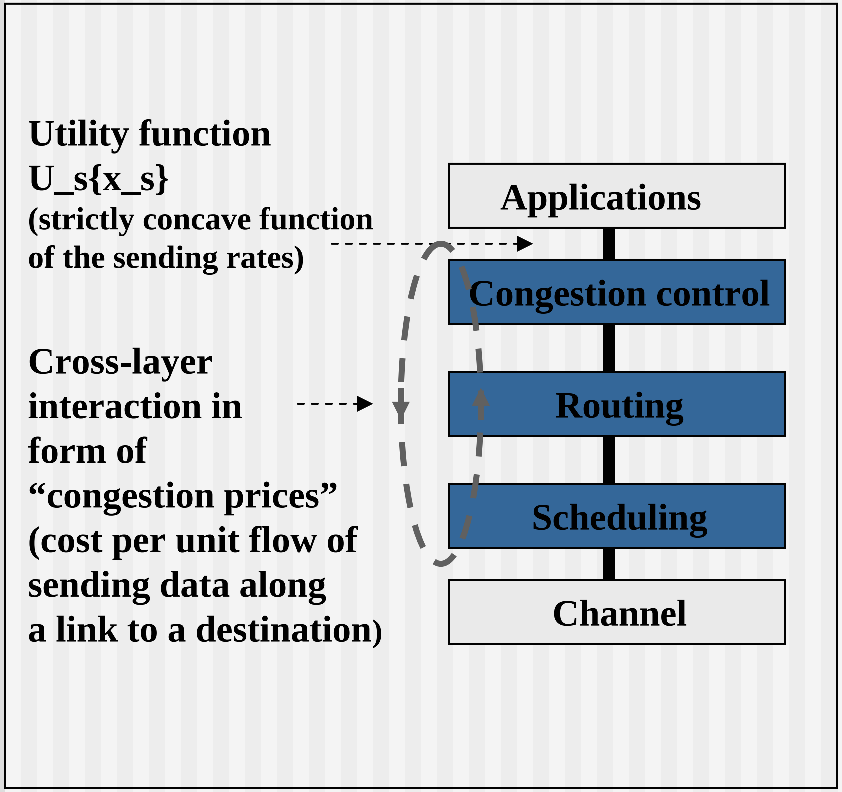


- MANET resource allocation formulated as global optimization problem
- Primal-dual decomposition generates a *set* of dual problems/algorithms/modules:
  - Local (except scheduling)
  - Tied together through congestion prices
- System Architecture traceable to theoretically provable optimality..

*Optimal Cross-Layer Congestion Control, Routing, and Scheduling Design in Ad Hoc Wireless Networks.* Lijun Chen, Steven H. Low, Mung Chiang[†], John C. Doyle (Caltech and [†]Princeton)

# Technology Experiments

- Why experiment with technology in GENI?
  - Not to test the technology itself, but to understand the relation of technology innovation to systems and architecture issues.
- Two examples:
  - Rapidly reconfigurable optical paths
    - change the relationship of routing, traffic engineering and topology management.
  - Wireless - particularly software defined wireless
    - raises issues of resource optimization, protocol structure, mobility management, locality management, coding and transport, handling unpredictable performance, …

# Advances in Network Design

- Cross layer design and optimization

- Routing
- Congestion control
- Addressing and naming architecture

- Service creation
- Autonomic management
- Self diagnosis
- ….

# Distributed Systems

- Distributed services and applications:
  - Robust content distribution
    - Naming, security, resilience
  - Management and sharing of personal information
  - Real time multi-media distribution
    - Multicast, data-driven dissemination, …
  - Network-embedded storage
  - Location mgt (human and object)
  - Identity mgt. (human and object)
    - Distributed name management

# What You Should Be Asking…

- GENI exists to serve the research community. We is us.

- What is the new idea I want to try?
- What will be needed in GENI to support it?
  - Technologies
  - Characteristics
  - Tools

# Where Might You Ask It?

- Many avenues for participation, or just feedback
  - These meetings
  - Mailing list, design discussions
  - Technical working groups

- The Research Coordination committee
  - Articulates the science case
    - To the customer communites
    - To other NSF communities
    - To congress and other sponsors
  - Make sure that GENI meets the needs of its customer communities.
    - We must anticipate that we will have to make choices here, for technical and resource reasons.

# GENI is not just infrastructure

- It is an opportunity to challenge both the community and NSF to think ambitiously about research goals, approaches and methods.

- GENI is not business as usual for CISE.
- FIND is not business as usual for NSF.
  - Collaboration and cooperation toward unified outcomes.
- Other NSF research programs may choose to restructure themselves.