


No Easy Answers

Viewing GENI in a Larger Context

John Wroclawski

USC/ISI





"Invention is the province of individuals, but its propagation through larger scale structures is dependent on their design, implementation and use"

- Andy Lippman - MIT Media Lab

- Andy was referring to technical structure...
...but social, economic, and business structures matter too; maybe more.
- Scope, framing and design of GENI directly impacts its ability to foster useful, practical and deployable solutions. This design must consider a larger context, not just a technical context.
- How're we doing?



Diversion: GENI design is not done



Step 1: Make the Problem Harder


Quote:

- “A central goal of GENI is to change the nature of networked and distributed systems design:
 - to apply intellectual rigor to the constraints of reality
 - to move from a fully empirical to a more rigorously understood design process”
- Transformational vs Evolutionary:
 - New designs vs new design paradigms

Example: Theoretically Derived Architectures



Electricity: 1800...



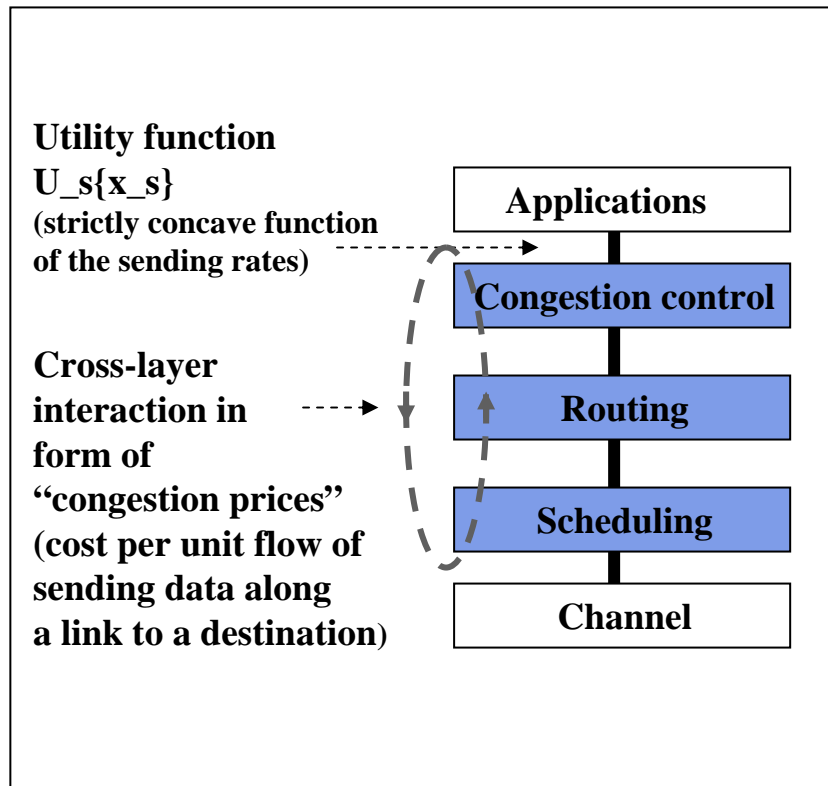
Electricity: Today...

$$\oint \vec{E} \cdot d\vec{A} = \frac{q}{\epsilon_0}$$
$$\oint \vec{B} \cdot d\vec{A} = 0$$
$$\oint \vec{E} \cdot d\vec{s} = -\frac{d\Phi_B}{dt}$$
$$\oint \vec{B} \cdot d\vec{s} = \mu_0 i + \frac{1}{c^2} \frac{\partial}{\partial t} \int \vec{E} \cdot d\vec{A}$$

(Network Architecture Today)

(Network Architecture ???)

Optimization-Based 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)
(and now a significant and growing body of publication..)



Paradigm Transformation

- Not impossible: already achieved in other engineering fields:
 - Electric power grids
 - Nuclear explosives
 - Bridges...
- GENI's goal: experimental validation of new design paradigms, as well as designs and system elements.
 - Transition of paradigm shifts??




- How do we compare paradigms and methodologies in ways that are convincing to the practitioner? What are the metrics?
- How to spread new paradigms through an industry?

Hard enough yet?

Step 2: Consider key aspects

(of the design, against the larger context)



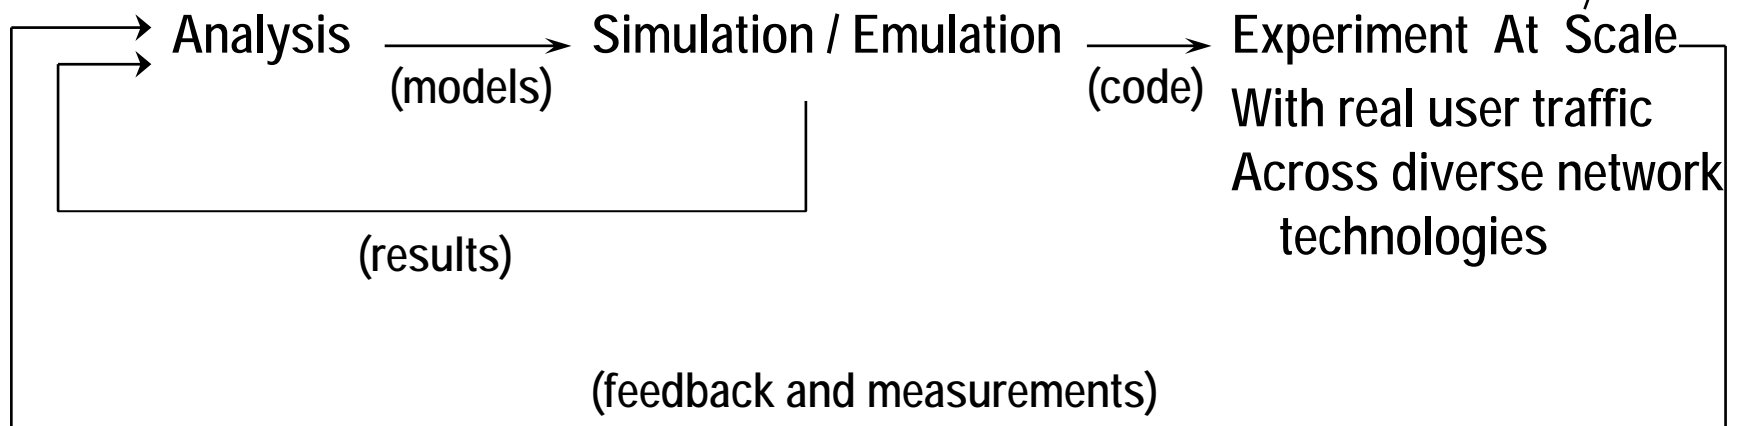
There are several. I discuss two:

- Focus
- Realism

Focus



- GENI documents discuss three purposes
 - Experimentation
 - Validation ("proof of concept")
 - Transition (the "Guru Gap")



One Facility, Three Purposes



- Have been somewhat blurred together, but not the same thing
 - (though similar in some dimensions...)
- Experiment:
 - controlled conditions: fix (freeze) some dimensions.
 - vary others, observe/measure results. - process of "science".
- Validation ("proof of concept")
 - emphasis on flexibility, rapid prototyping methodologies.
 - some exposure to potential early adopter users/customers.
- Transition
 - not necessarily a technical issue.
 - exposure to non early adopter users (the famous chasm).
 - heightened requirement for reliability, dependability, security, maintainability, ..
 - may require significant scaling.
 - may require more extensive economic/commercial model.



- Multiplicity of purposes => potential tradeoffs
 - Can (a single) GENI really do all of this?
 - Are there alternatives - some additional structure - that might do it better
- What might a more structured GENI look like?
 - At the facility level? At the slice level?
 - Tiered/recursive? - little GENI big GENI
 - Structurally driven federation - many facilities with different foci? Parallel evolution?
- Keys:
 - Ensure that each focus is understood/described independently
 - Ensure that necessary conditions/inputs are available for each
 - Design GENI to share common dimensions, isolate others

Realism



- GENI documents: key goal (is) supporting experimental research under realistic conditions.
- What does this mean?
- What are the dimensions of realism?
- How does the answer to this question affect usefulness of GENI to academia and industry?



Realism

Some common dimensions:

- "Real users"
- Real traffic and traffic matrix
- "Realistic topology"

(obviously, a certain amount of judgment for any given experiment..)



But wait..

- Relative cost of components *in production*
 - (access vs core, wireless vs wired/optical, etc)
- Relative Moore's Law
- Network management overhead/requirements in production situations
 - common case (operational)
 - exception case (fault management)
- Legal environment
- Regulatory environment
- Marketing environment (somewhat doable?)
- Industry structure & relationships
-



Some of these dimensions not considered in
the current design :)

ISP Thought Experiment

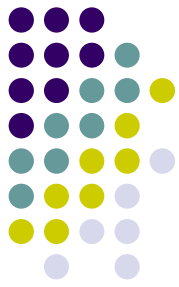


- Why do ISP's exist?
 - Repository of capital
 - Repository of specialized skills
 - Network management
 - Business analysis / legal / regulatory
 - Construction - facilities
 - ...




What if questions

- Wireless mesh networks:
 - Dramatically change the requirement for capital needed to build out urban access infrastructure
 - May or may not change the requirement for repository of management, facilities skills.
- Johnny Depp: “That’s Interesting...”



- Q1: How might wireless mesh networks change the structure of the industry?
 - Who will make money? Where should you invest development resources? What is the value chain?
 - Corollary:
 - Are structural changes required? Or merely catalyzed?
- Q2: same question: WM -> net virtualization
 - From the technical perspective, a very hot topic.
 - From a larger socio-business-economic perspective, (nearly) untouched..
- Q3: Are these GENI questions at all? Should they be? How might GENI address questions like this?



- (Q3: Are these GENI questions at all? Should they be? How might GENI address questions like this?)
- A1: Nope. Too hard, and a distraction.
- A2: Try.
 - Incorporate into the GENI toolbox
 - Skeleton models for components, industry structures, etc
 - Analytic tools (Real Options, Game Theoretic toolkits*, ...)
 - Ability to provide data and run what-if scenarios
 - Broaden the GENI design process (or a corollary effort outside of NSF-GENI) to more aggressively reach economists, social scientists, and legal/regulatory experts as candidate research team members

*Skepticism duly noted and somewhat agreed with

