Workshop on Network Science and Network Design

July 29-30, 2008 at USC/ISI

Purpose…

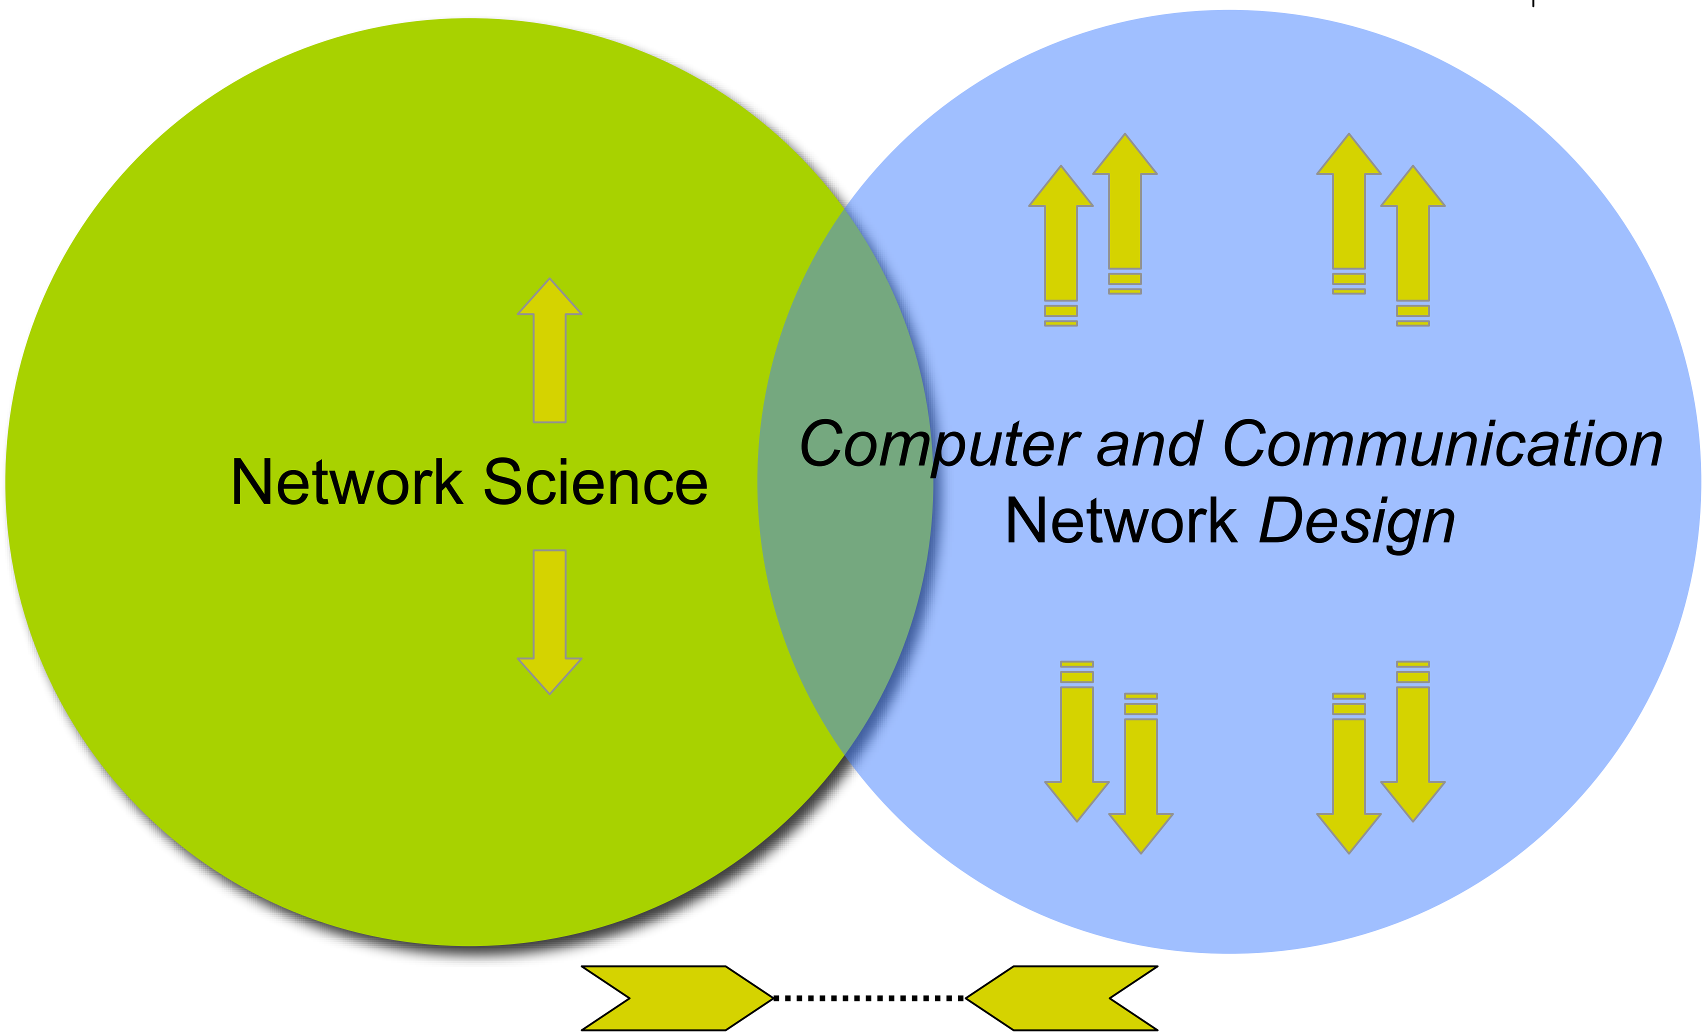
Network Science
Computer and Communication
Network *Design*

People & Path

- David Alderson
  Naval Postgraduate School
- David Clark
  MIT CSAIL
- Heidi Picher Dempsey
  GENI Program Office
- John Doyle
  Caltech
- Darleen Fisher
  NSF
- Fan Chung Graham
  UC San Diego
- Suzanne Iacono
  NSF
- Ali Jadbabaie
  Penn
- Will Leland
  Telcordia Research
- Dmitri Krioukov
  CAIDA
- R. Srikant
  U of Illinois
- Walter Willinger
  AT&T Research
- John Wroclawski
  USC ISI
- Ellen Zegura
  Georgia Tech
- Ty Znati
  NSF

Views of Network Science

- Different people use this phrase in different ways
- An established community with its own culture and perspective
- Views:
  - “any theory that has to do with networks”
  - “power laws and scale free graphs”
  - Search for common abstractions, metrics, tools *across network domains*

Evolution…

- Search for common abstractions, metrics, tools *across network domains*
  - Powerful. Tricky.

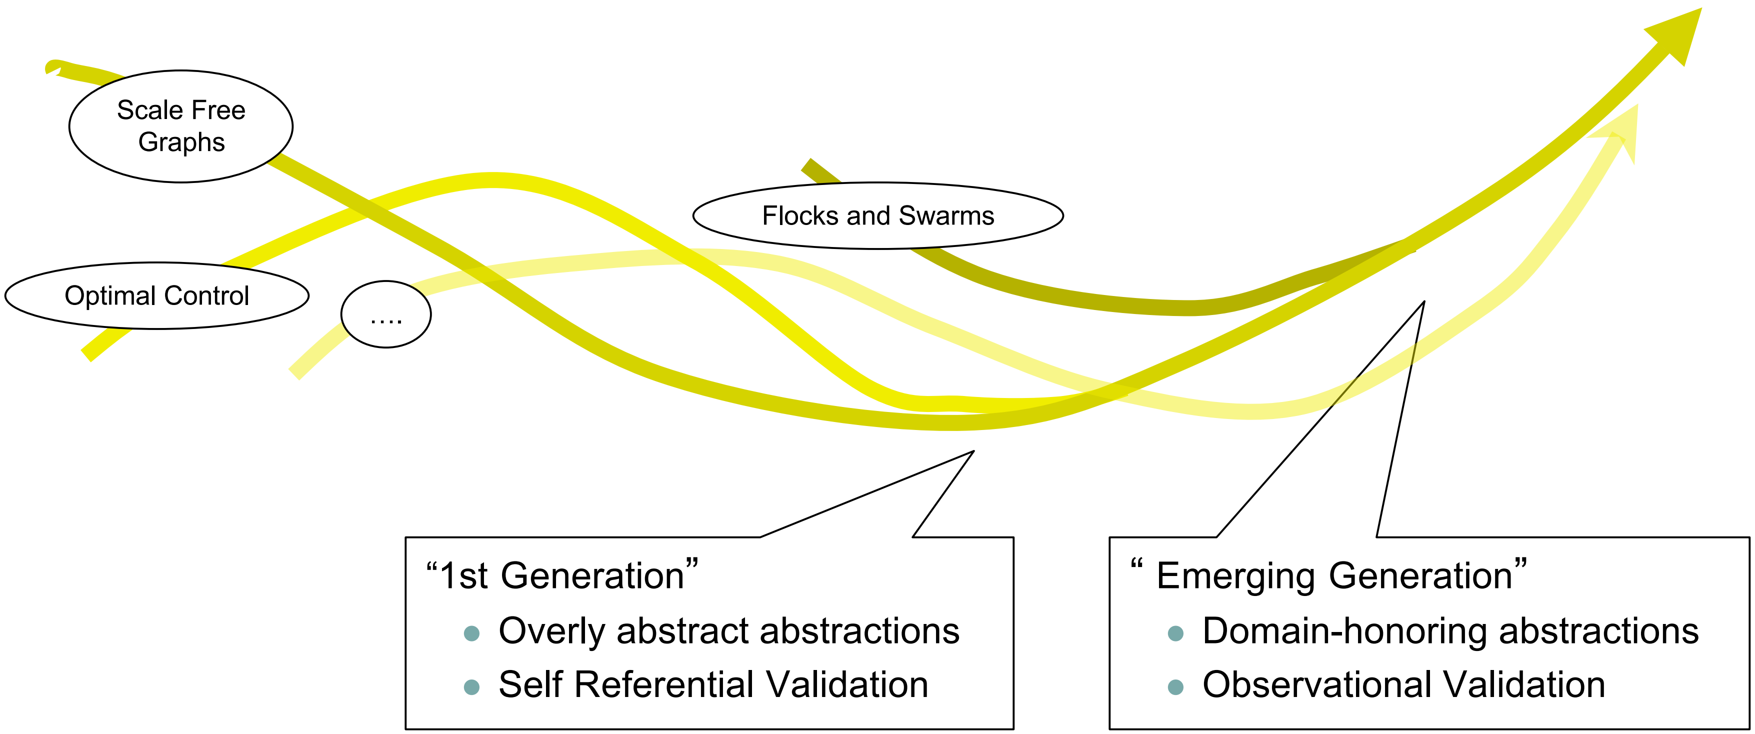

The structure of scientific explanation

- Different sciences use levels unevenly and in very different ways.

Levels

1. Verbal
2. Data & stats
3. Modeling & sim
4. Analysis
5. Design & synth

- Network ***science*** has historically focused on data and statistics.
- Network ***design*** demands a fundamental rethinking / transition, particularly ***proofs*** in ***analysis and synthesis***.

Evolution of
Theory and the Internet

Goals

- Abstraction (common concepts across fields)
- Rigor (& math structure)

Issues

- Statics (topology, structure)
- Dynamics (location, propagation)
- Robustness (& security)

Levels of understanding

1. Verbal (& cartoons)
2. Data & statistics (Experiments & measurements)
3. Modeling & simulation
4. Analysis
5. Design & synthesis

Good news:

- Spectacular progress

Bad news:

- Persistent errors and confusion
- Potentially insurmountable obstacles?

This slide run courtesy John Doyle, slightly edited

"The Matrix"

- subfields of networking, and progress therein..

	Traffic	Topology
Verbal		
Data/stat		
Mod/sim		
Analysis		
Synthesis		

ARCHITECTURE

C&D	Layering	???

A success story
Traffic (1993-2000)

	Traffic
Verbal	
Data/stat	
Mod/sim	
Analysis	
Synthesis	

- Heavy tails (HT) in net traffic???
- Careful measurements
- Appropriate statistics
- Connecting traffic to application behavior

  HT files ⇒ HT traffic
- “optimal” web layout

A lesson learned
Topology (1999 - Present)

	Traffic	**Topology**
Verbal		
Data/stat		
Mod/sim		
Analysis		
Synthesis		

Control, Dynamics, and Architecture:
Exciting recent progress

	Traffic	Topology
Verbal		
Data/stat		
Mod/sim		
Analysis		
Synthesis		

ARCHITECTURE

C&D	Layering	???

Three Research Examples

Extending a Theory

New Columns in the Matrix

Design by Constraint

Example: Extending the Theory
Theoretically Derived Architectures

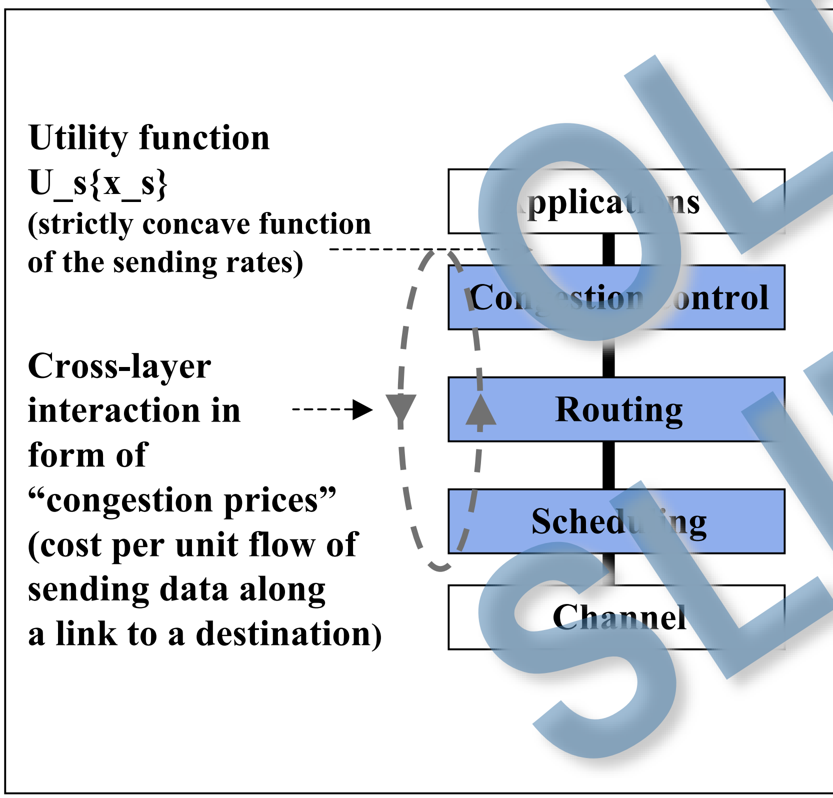

- Network resource allocation formulated as global optimization problem
- Primal-dual decomposition generates a *set* of dual problems/algorithms/modules:
  - Local (except scheduling)
  - Tied together through congestion prices
- System <u>Architecture</u> traceable to theoretically provable optimality..

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)

Example: Extending the Theory
New Challenge: Delay

- Previous work structured around *flow rate*
- *Delay* is the critical issue for many new applications:
  - Cyber Physical Systems (Networked Control)
  - Games, Interactive Communication, etc

- Approach: (attempt to) apply a tested methodology..
  - Enhance modeling to capture new effects (OK)
  - Identify and add new constraints to optimization problem (~OK)
  - Extend theory to operate in the presence of new constraints (So far, hard..)

- *Key result* if successful:
  Theoretically derived architecture for delay-sensitive networks

History: Continual Advance *through* Similar Methodology

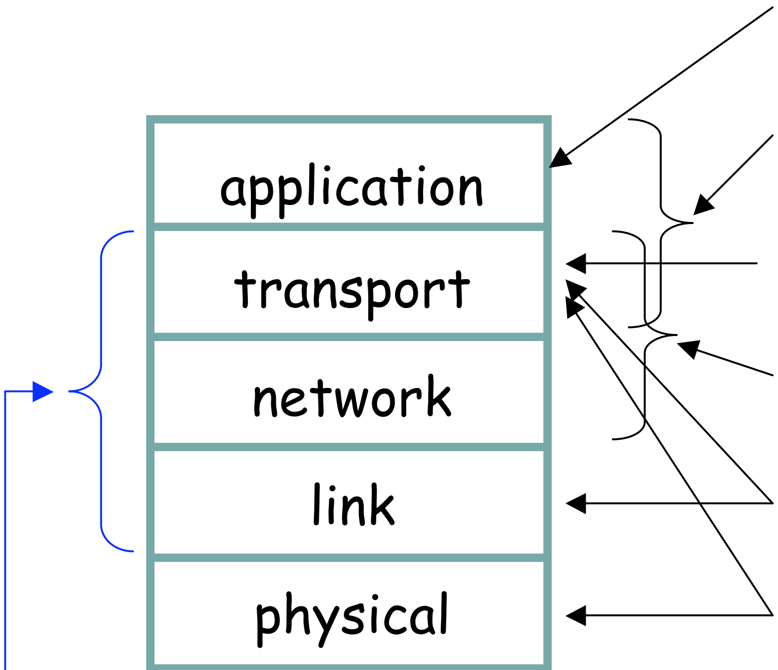

Web layer: Zhu, Yu, Doyle ’01

HTTP/TCP: Chang, Liu ’04

TCP: Kelly, Maulloo, Tan ’98, ……

TCP/IP: Wang et al ’05, ……

TCP/MAC: Chen et al ’05, ……

TCP/power control: Xiao et al ’01, Chiang ’04, ……

Rate control/routing/scheduling: Eryilmax et al ’05, Lin et al ’05, Neely, et al ’05, Stolyar ’05

Integrating network coding w/above: (Chen et al ’07, Cui et al ’07, …)

Detailed Survey: Proc. of IEEE, 2007

Example: New Columns in the Matrix "Security" in a theoretical framework?

Caveat: entire slide is insane speculation

- Challenge:
  broaden theoretical frameworks to include additional design elements

- Key issues:
  - Metrics
  - Relatable Metrics

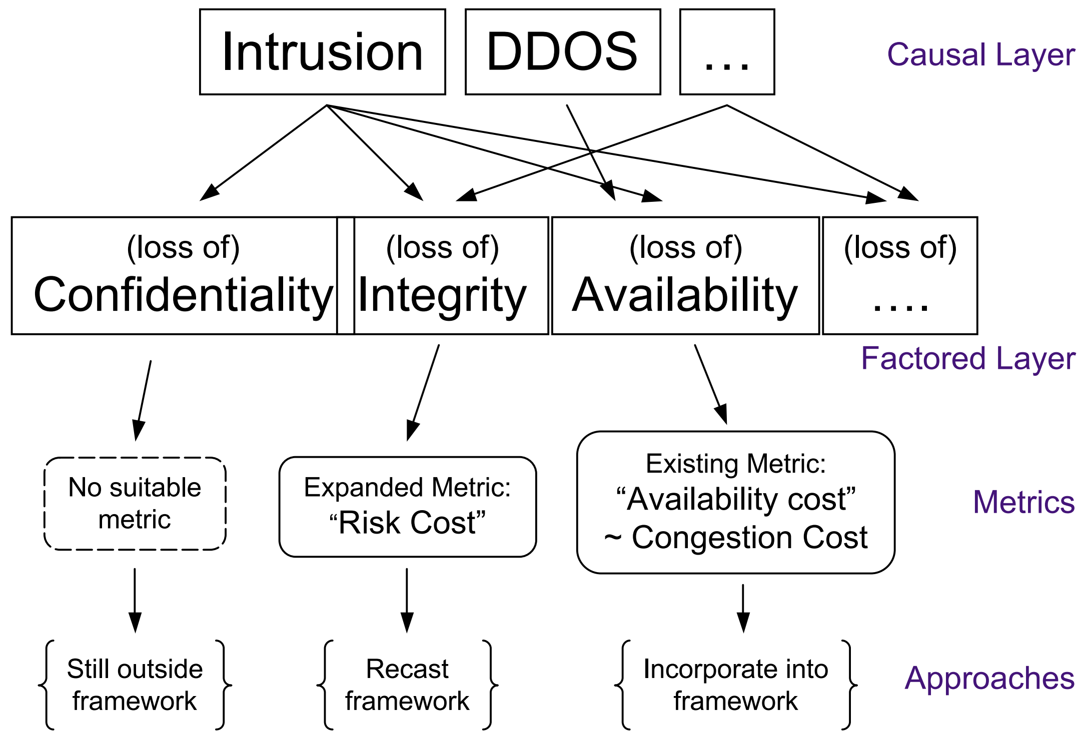

Example: Global Results from Local Actions
"Design by Constraints"

Two "Internet Topologies"; same power law parameters..

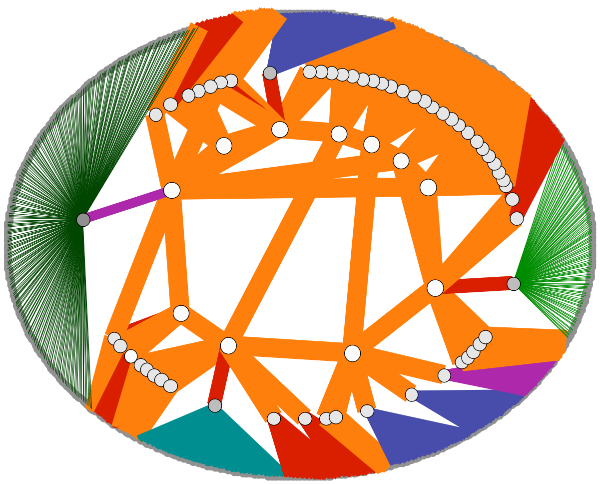

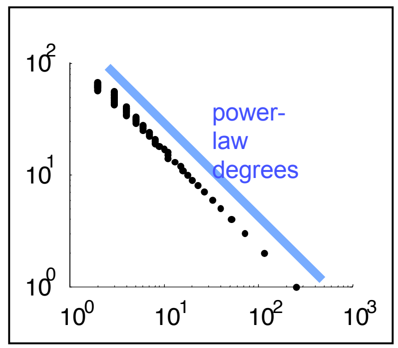

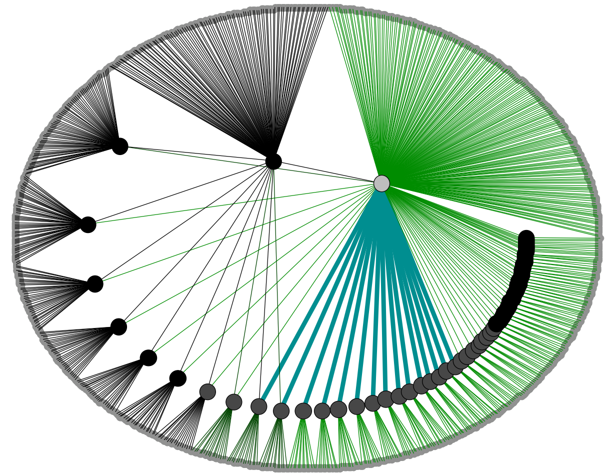

Low degree mesh-like core

- *High* performance and robustness
- Efficient, economic
- From "random" generator, *low* probability, but
- *Like* real Internet

High degree hub-like core

- *Poor* performance and robustness
- Wasteful, expensive
- From "random" generator, *high* probability, but
- *Unlike* real Internet

What's going on?

- This is *surprising* to many in network science
- This is *not surprising* to most Internet engineers
- What's going on?

Abilene Backbone Physical Connectivity

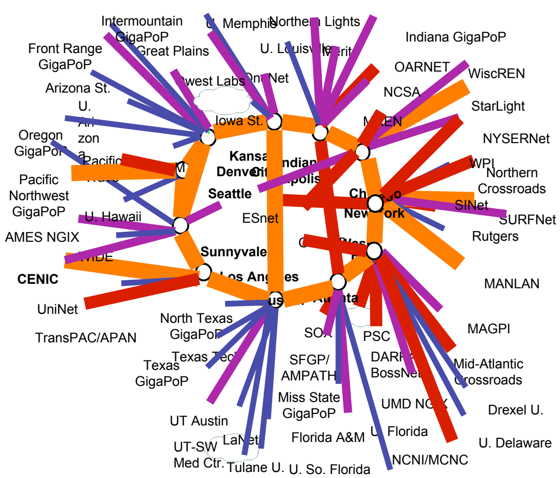

Start with an *engineered* backbone…

Add gateway routers and end systems *consistently with technological constraints* on these routers and systems…

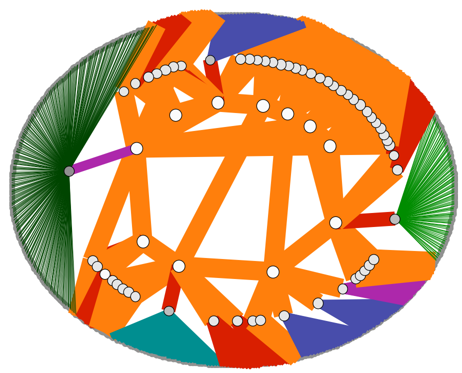

Get topology [synthesized or real] with high throughput, efficiency, economy

Design by constraint

- The desirable topology is due to both
  - Classical engineering
  - *Local constraints* shaping *global results*
- To be fair, perhaps somewhat by accident..

- The key question: can we do it <u>on</u> <u>purpose</u>?
  - Design, not of the complete system, but of components from which systems with desired properties will come forth
  - Formalization of methods for this class of design

Validation

"I argue that power law research must move from focusing on observation, interpretation, and modeling of power law behavior to instead considering the challenging problems of validation of models and control of systems"
- Michael Mitzenmacher*

- *Validation* is essential to the progress of science and engineering
- We need some attention here…
- Three types:
  - "Self Referential Validation"
  - Observational Validation
  - Generative (design-based..) Validation

**Editorial: The Future of Power Law Research*.
M. Mitzenmacher, Internet Mathematics 2(4), 2006

Self-Referential Validation

- Identify a phenomenon in/of the original artifact
- Develop a mathematical model that captures that phenomenon..
- Note similarity.

- Key question: what is being validated here?

Observational Validation

- "Classic Science"
  - Model artifact based on observation of phenomena
  - Use model to predict *different* correlating/supporting phenomena
  - Observe artifact to validate model
- Key questions
  - Correlation between primary and validating phenomena
  - Observational platform capabilities (Internet)
    - Today, often focused on observing primary phenomena
    - Partial information and other observational problems

Generative Validation

- "Build it and see/study what happens"
  - For computer systems, often the most convincing approach
  - For standard engineering problems, no "fundamental" difficulty
- Key question: Type 3 problems*
  - Validating evolution over time
  - Validating the results of others' actions..
- The heart of our challenge, yet the hardest of all

*("global results from local action", "design by constraint", etc..)

Generative Validation: Ideas

- World models, not system models
  - Economics
  - Adaptive/intelligent user models
  - External event models
  - …
- Artificial environments based on worst case analysis
- World simulation / system emulation
- …

Workshop on Network Science and Network Design

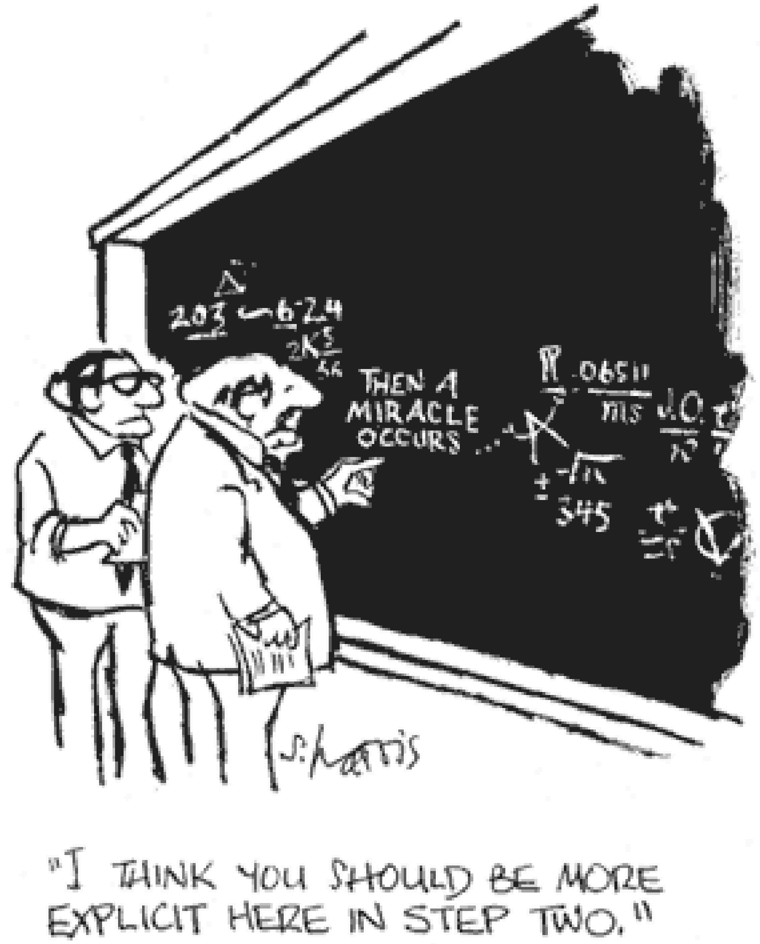

"Just a little bit further to go.."