Diverse Network Infrastructure for Resilience and Rapid Recovery from Large-Scale Disasters

James P.G. Sterbenz*†
Egemen Çetinkaya*, Abdul Jabbar¶*, Justin P. Rohrer§*
David Hutchison†, Paul Smith◊†, Marcus Schöller‡†
Deep Medhi, Jiannong Cao, Jinyao Yan

*Department of Electrical Engineering & Computer Science
Information Technology & Telecommunications Research Center
The University of Kansas
†School of Computing and Communications, Infolab 21
Lancaster University, UK

jpgs@ittc.ku.edu *jpgs@comp.lancs.ac.uk*
http://www.ittc.ku.edu/~jpgs
http://wiki.ittc.ku.edu/resilinets

Evaluation of Network Resilience

Abstract

The Internet is a critical infrastructure on which we depend, and thus it is essential that it be resilient such that it continues to provide service in the face of various challenges, including attack and large-scale disasters such as earthquakes, hurricanes, tsunami, and coronal mass ejections. A major aspect of our previous and current work on achieving resilience has centered on providing diversity in the network such that when part of the network fails, alternatives will be available to continue operation. This includes heterogeneity and diversity in mechanism (for example wired and wireless), rich topology interconnection, and structural diversity of the network graph such that paths can be constructed that do not share fate when network components fail. We have developed a set of analytical and simulation techniques and tools to generate network topologies, and to analyse the resilience of real and synthetic network graphs. We have also made our data and topology viewer publicly available at http://www.ittc.ku.edu/resilinets/maps. A key aspect of this ongoing work is the multilevel nature of the analysis, in particular, attacks against the physical infrastructure must be modelled on the physical layer graph (fiber interconnection for the typical backbone) but its effects analysed on the IP network layer graph overlay. Under new funding from NSF NeTS (in collaboration with Deep Medhi at UMKC), we are exploring geographic diversity and its impact on traffic load, such that networks can be designed to survive large-scale disaster of a given scope. For example, an application should be able to specify: give me three multipath routes over which communication can be erasure coded such that the paths are no closer than 100 km (except at the source and destination). This example defends against a disaster with a diameter of less than 100 km in diameter. While we work to understand how to generate networks with desired graph-theoretical, diversity, and resilience properties, the reality is that even if adopted, there will be cases where disasters will partition the network, either because the area is greater than anticipated, or cost constraints have not permitted the deployment of sufficiently resilient infrastructure.Thus, we are beginning research on how to optimally and rapidly deploy infrastructure after a disaster, in particular, to restore services outside the disaster area, to rapidly deploy assets to permit assessment of the damage to the environment and network, and to rapidly and optimally deploy infrastructure to restore critical network infrastructure to the affected area. This is joint work with Chinese instutions (Jiannong Cao at Hong Kong Poly and Jinyao Yan at CUC Beijing) as a result of our participation in the US NSF / China NSFC Workshop on Environmental Monitoring or Public Health and Disaster Recovery.

Resilient Networks
Motivation

- Increasing reliance on network infrastructure

⇒ Increasingly severe consequences of disruption

⇒ Increasing attractiveness as target from bad guys

  - recreational and professional crackers
  - industrial espionage and sabotage
  - terrorists and information warfare

Resilient Networks
Definition

- Resilience [*ComNet* 2010]
  - provide and maintain acceptable service
  - in the face of faults and challenges to normal operation
- Challenges [DRCN 2013]
  - unintentional misconfiguration or operational mistakes
  - large scale disasters (natural and human-caused)
  - malicious attacks from intelligent adversaries
  - environmental challenges
  - unusual but legitimate traffic
  - service failure at a lower level

ResiliNets Strategy
D^2R^2 + DR

- Two phase strategy for resilience
- Real time control loop: D^2R^2
  - real-time with respect to network operation
  - many simultaneous independent loops
- Background loop: DR
  - out-of-band analysis of the reaction to adverse events
  - increase future resilience of system

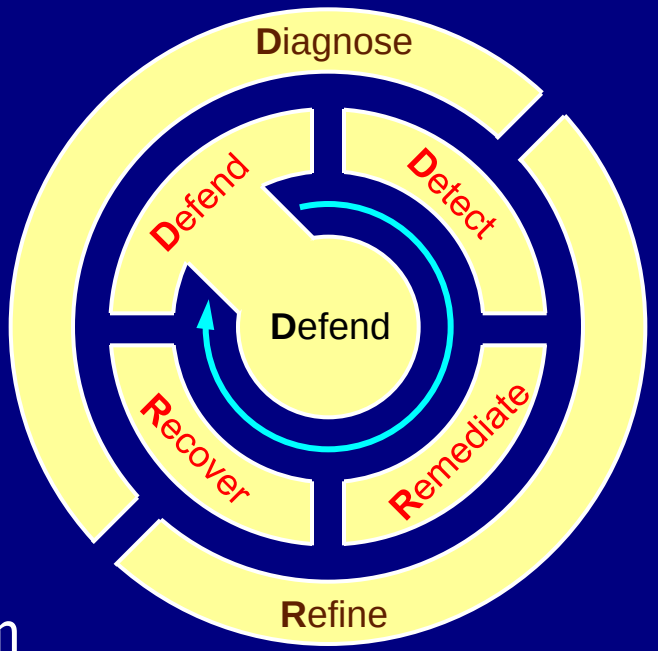

[wiki2006, EU ANA, EU ResumeNet, NSF PoMo, *ComNet* 2010]

ResiliNets Strategy
D^2R^2 + DR: Defend (Passive)

S1a. Defend against challenges to normal operation

- Reduce the probability of a fault leading to a failure
- Reduces the impact of an adverse event
- Disaster tolerant network
  - spatially diverse redundant paths
  - medium diverse paths
  - robust power alternatives
  - weather- and rad-hard components

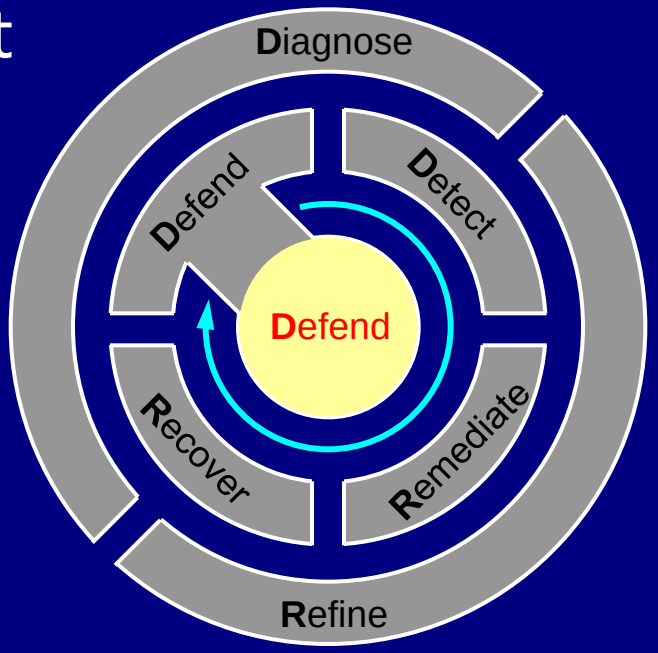

ResiliNets Strategy
D^2R^2 + DR: Defend (Active)

S1b. Defend against challenges to normal operation

- Reduce the probability of a fault leading to a failure
- Reduces the impact of an adverse event
- Disaster tolerant network
  - active disaster resistance with interdependent infrastructure
  - flash-crowd tolerance
  - filtering traffic for known attack signatures

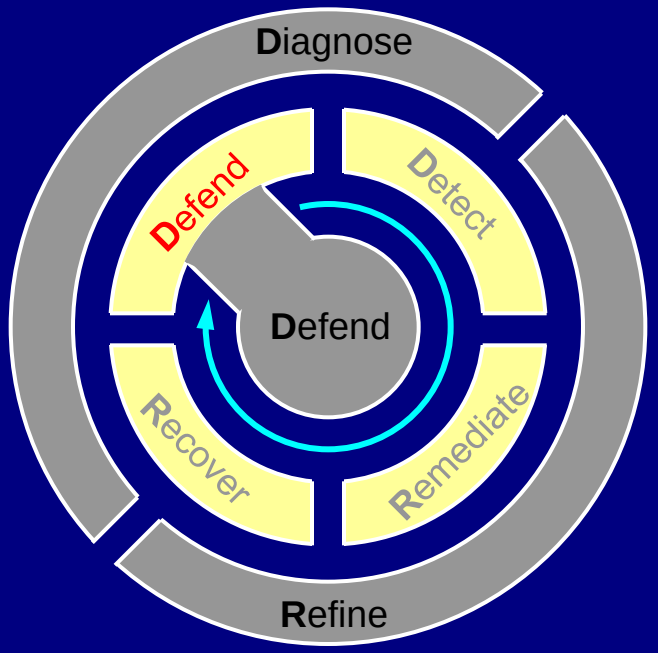

ResiliNets Strategy
D^2R^2 + DR: Detect

S2. Detect when an adverse event or condition occurs

- Determine when defenses
  - have failed and remediation needs to occur
  - need to be strengthened
- Disaster tolerant network
  - detection of behavioural anomaly
    - traffic load or pattern
  - alarms when infrastructure fails
  - sensors to assess damage
    - safety to first responders and net engineers

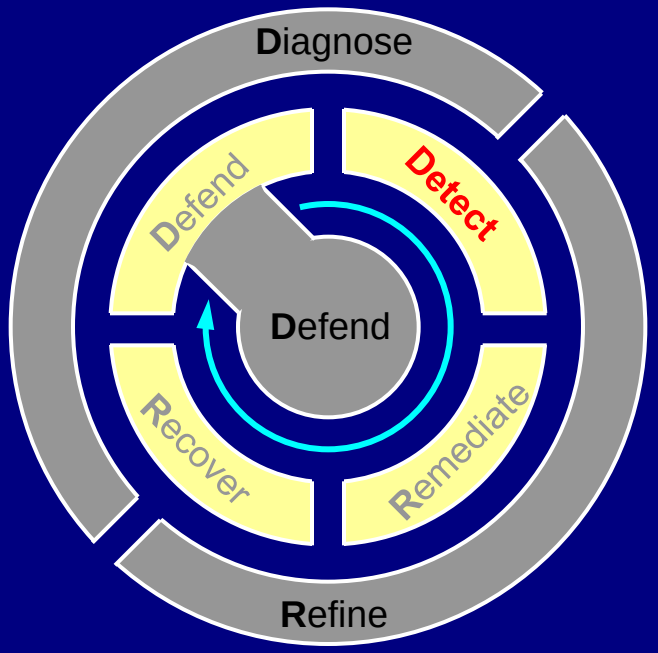

ResiliNets Strategy
D^2R^2 + DR: Remediate

S3. Remediate during adverse condition

- Do the best possible
  - after/during adverse event/condition
- Corrective action at all levels
  - graceful degradation
- Disaster tolerant network
  - reroute network traffic around
  - rapid deployment of temp net

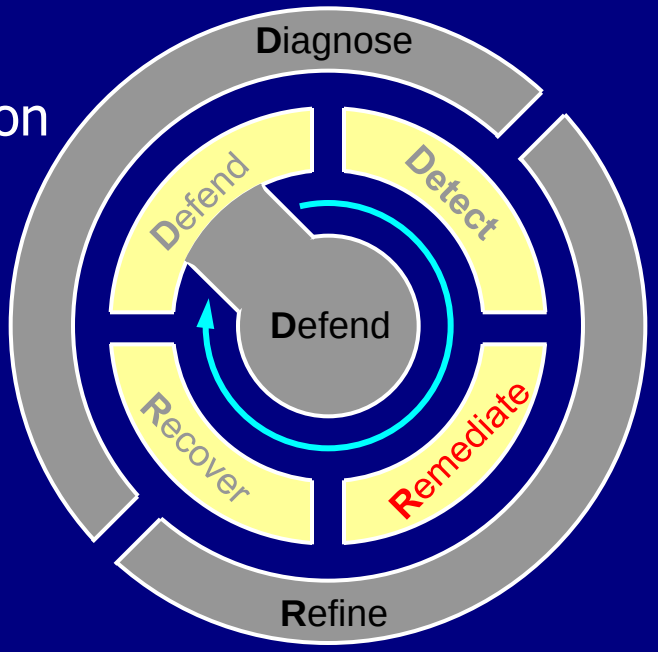

ResiliNets Strategy
D^2R^2 + DR: Recover

S4. Recover to normal operations

- Return to original state once adverse condition over
  - redeploy infrastructure
  - restore normal control and management
- Disaster tolerant network
  - restore original network routing
  - replacement of infrastructure
    - subject to…

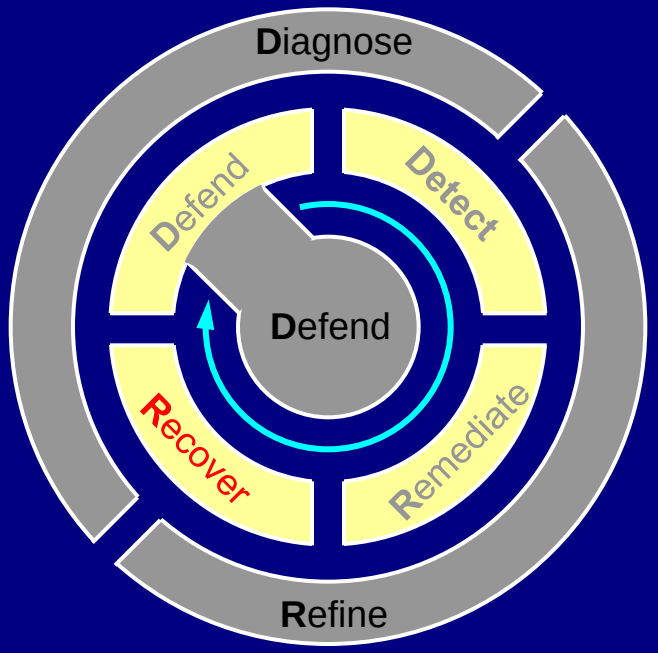

ResiliNets Strategy
D^2R^2 + DR: Diagnose

S5. Diagnose fault that lead to error or failure

- Root cause analysis to discover design flaws
  - faults not directly detectable
- Disaster tolerant network
  - analyse disaster
  - network response
  - root cause analysis

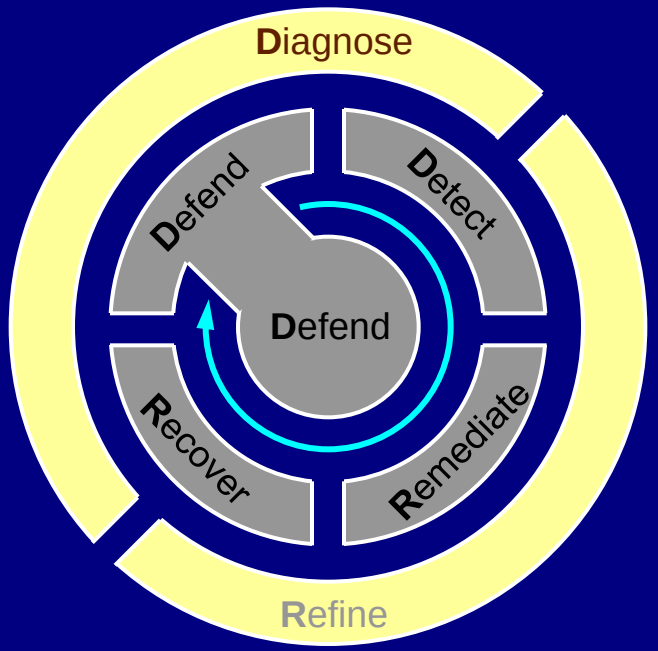

ResiliNets Strategy

D^2R^2 + DR: Refine

S6. Refine behaviour for the future

- Learn from past D^2R^2 cycles
  - better defense, detection, remediation *next time*
- Disaster tolerant network
  - modify threat models
  - modify challenge matrix
  - enrich network topology
  - redesign protocols
  - strengthen resistance to cascading and interdependent failures

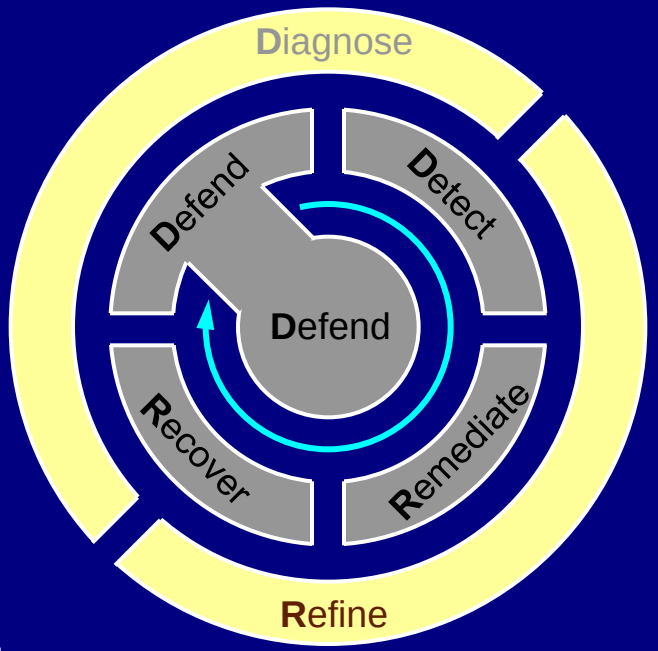

Resilience Principles
Redundancy, Diversity, Heterogeneity

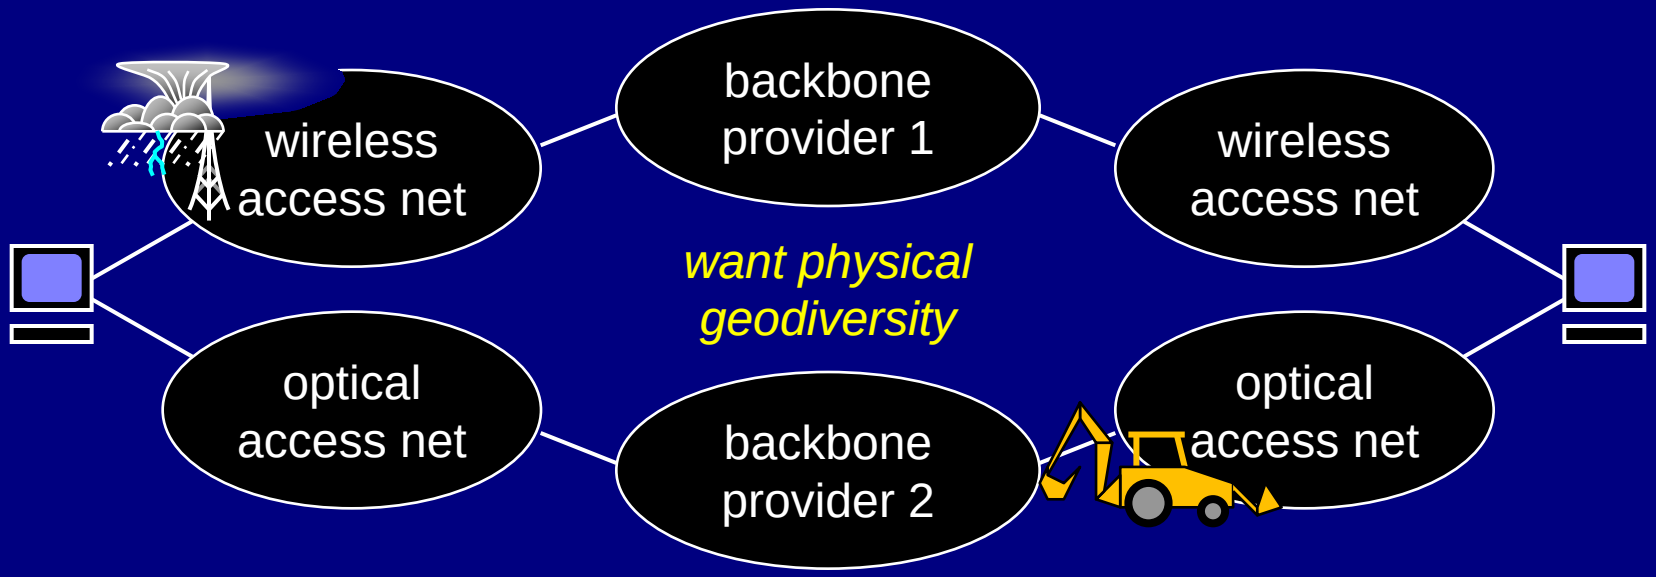

- Realm path choices explicitly available to end user
  - spreading (e.g. erasure coding) or hot standby
  - service tradeoffs: optical when available, fail-over to wireless
  - cheapest path under dynamic pricing

ResiliNets State Space
Operational Resilience

- Operational resilience
  - minimal degradation
  - in the face of challenges
- Resilience state
  - remains in normal operation

[ICNP 2006, COMSNETS 2011, *TSJ* 2011]

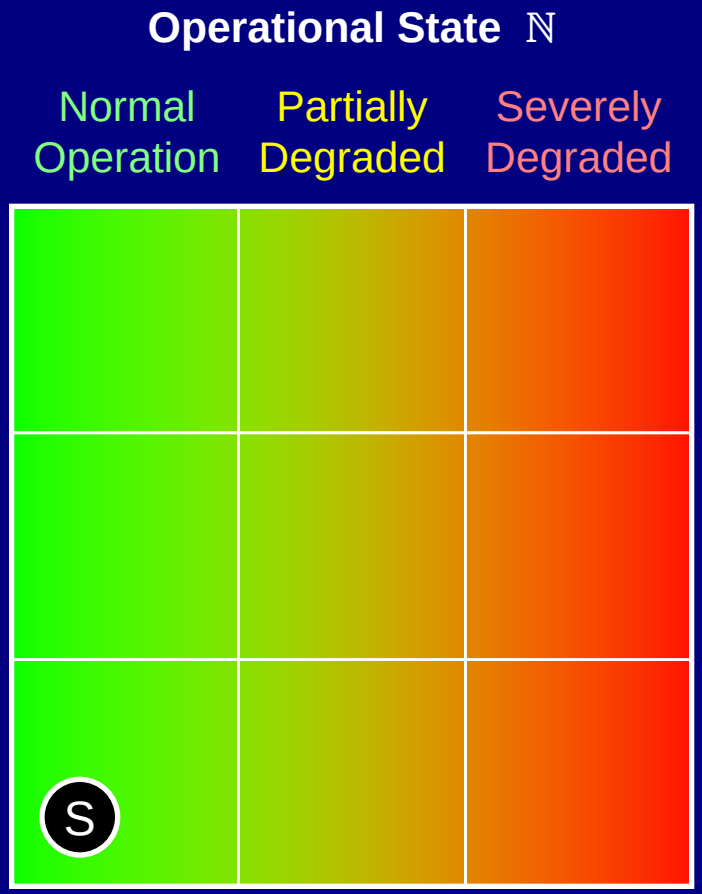

ResiliNets State Space
Service Resilience

- Service resilience
  - acceptable service
  - in the face of degraded operation
- Resilience state
  - remains in acceptable service

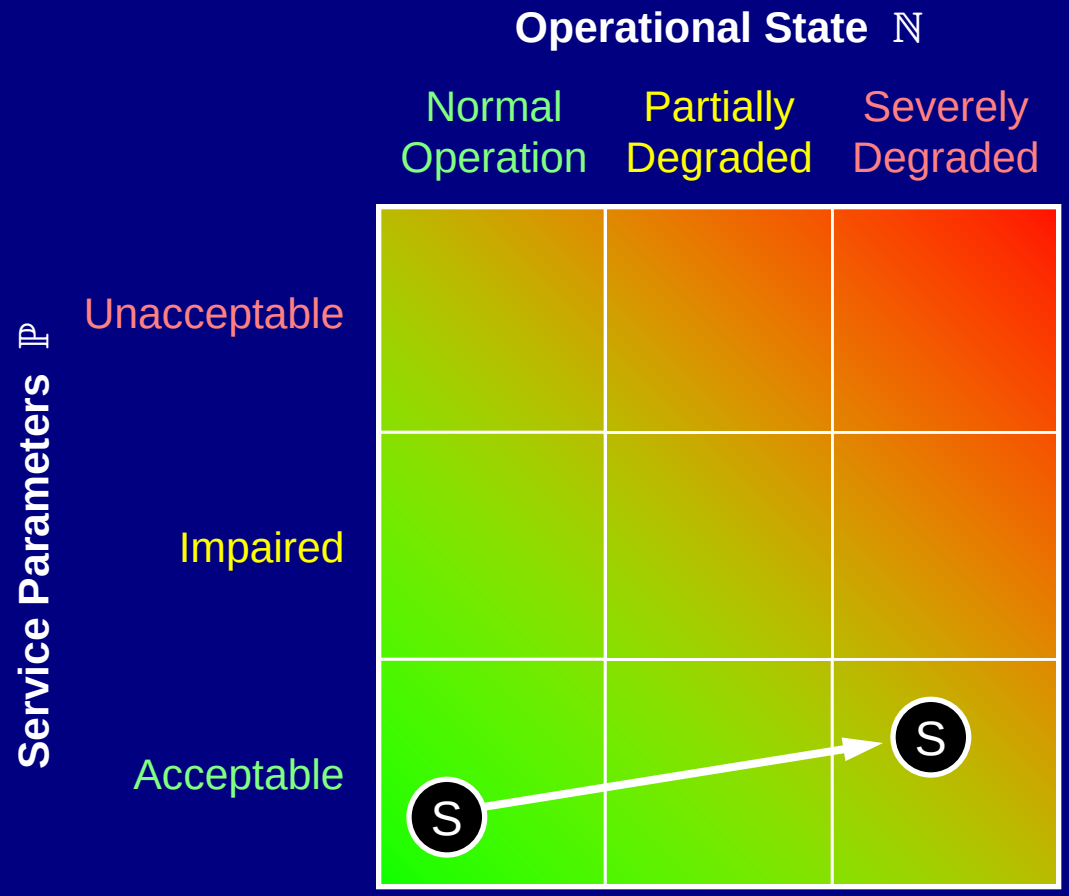

ResiliNets State Space
Quantification of Resilience

- Resilience
  - $\mathbb{R}$ = area under trajectory
  - for particular scenario
  - resilience $\Re$ over all scenarios
- Types of analysis
  - static [Jabbar 2010]
  - temporal
  - reflective

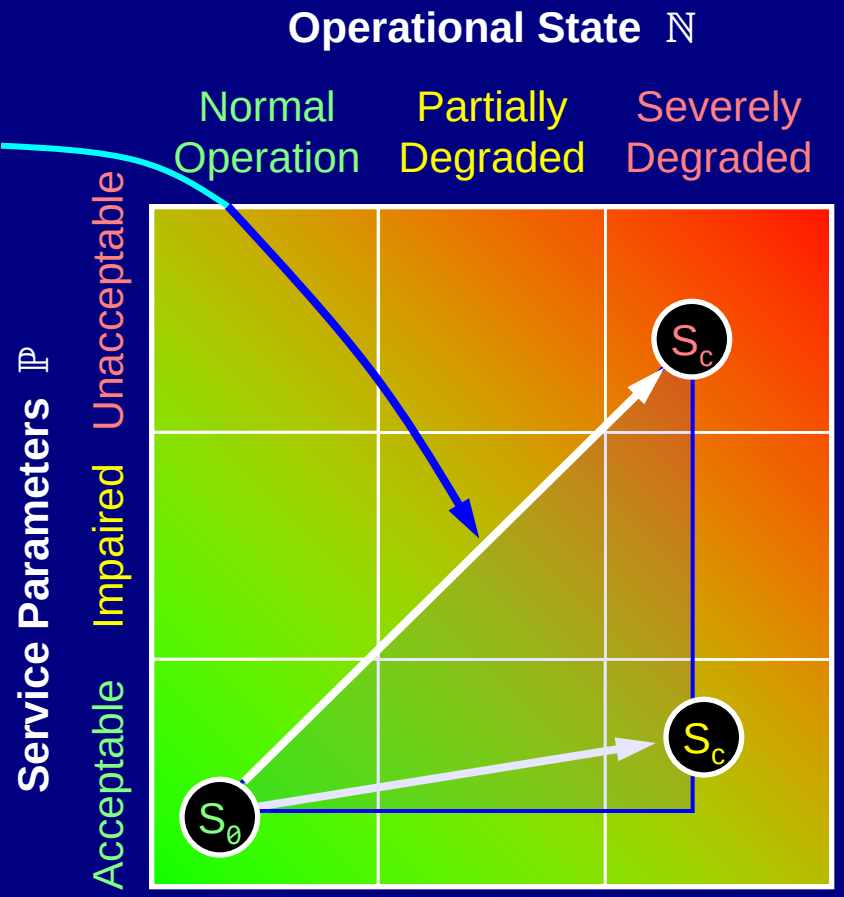

ResiliNets State Space
Analysis Alternatives

- Resilience
  - 1–$\mathbb{R}$ (area under trajectory)
  - for particular scenario
  - resilience $\mathfrak{R}$ over all scenarios and levels
- Types of analysis
  - static [Jabbar 2010]
  - temporal: weight by time
  - reflective: compare refinement alternatives

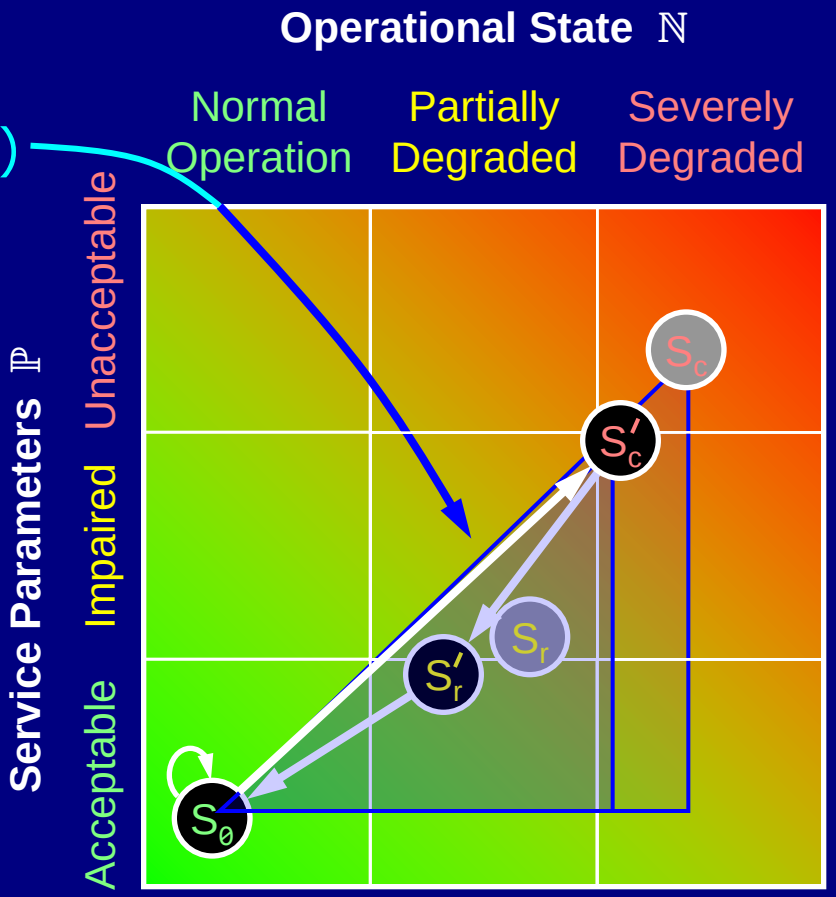

State Space Analysis
Network Topology Example

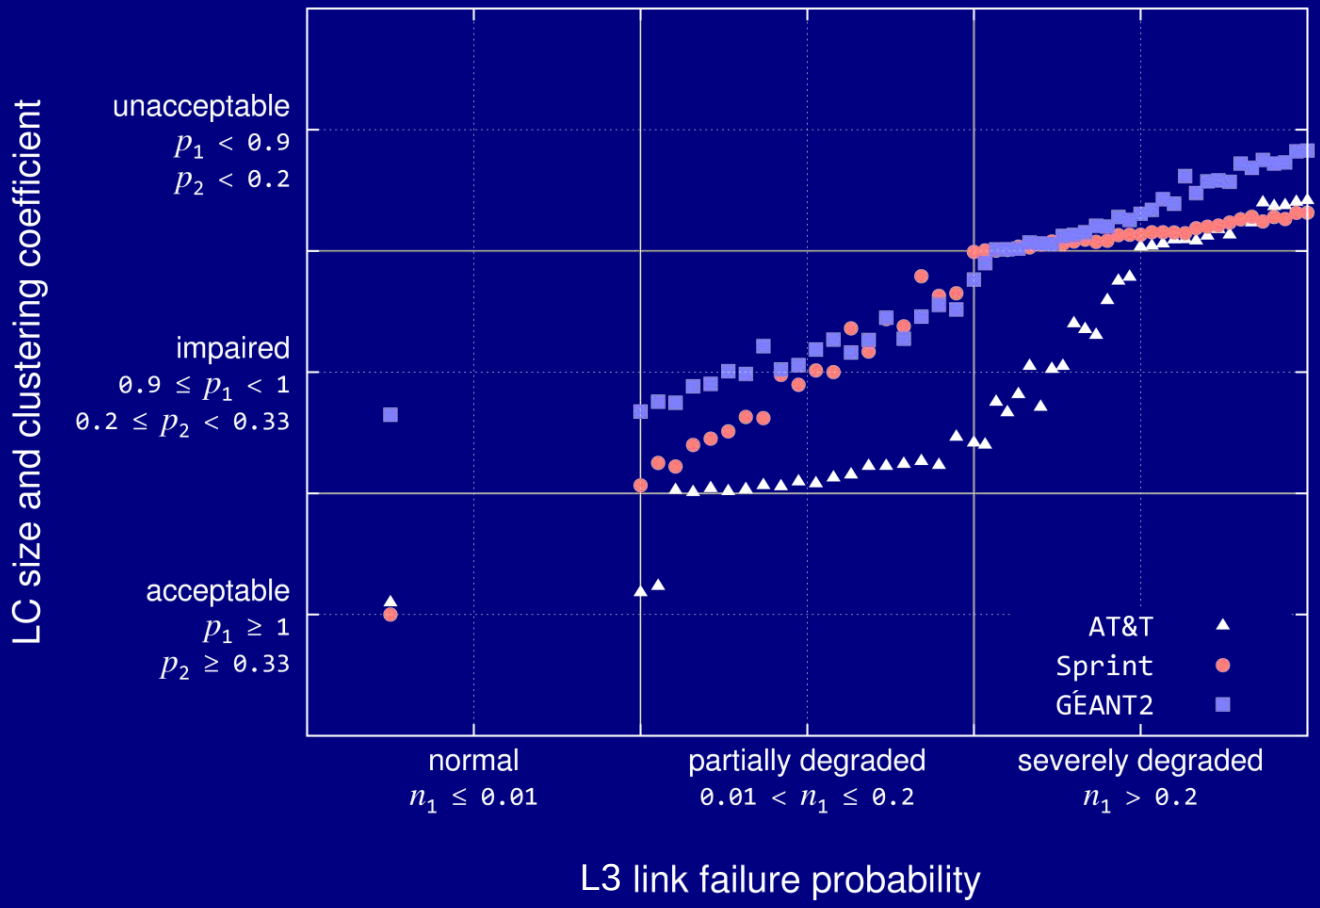

ResiliNets State Space

Resilience of Disaster Recovery

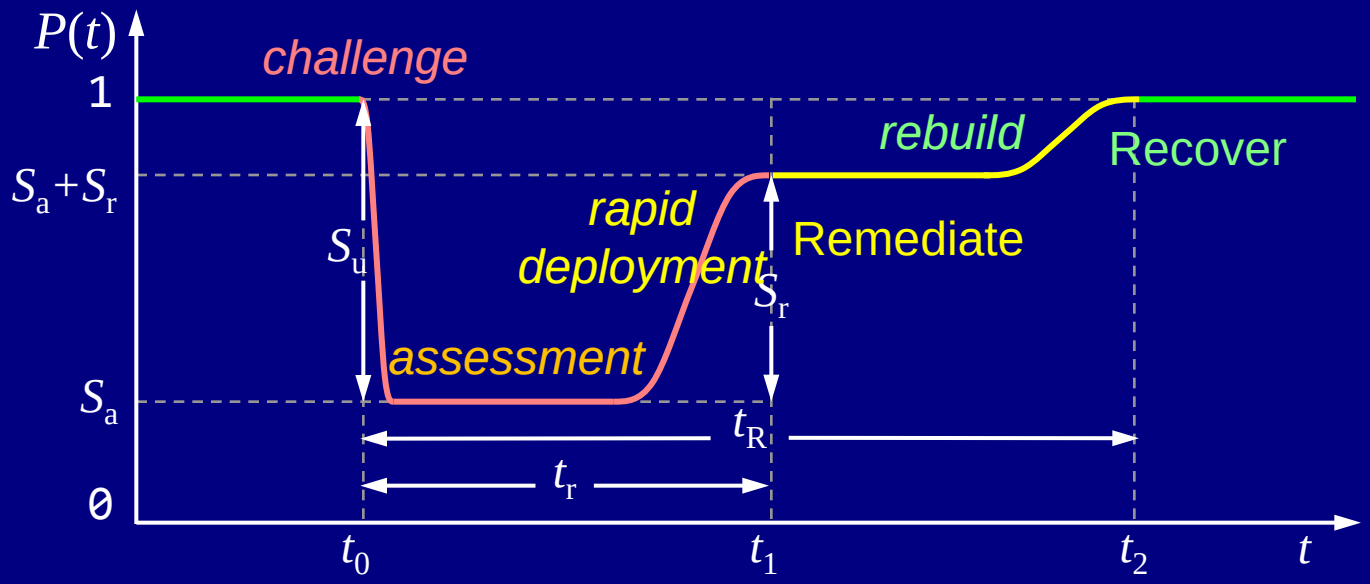

- Time to remediate t_r; to recover t_R
  - temporal weights applied to $\mathbb{R}(t_0, t_R)$
- After refinement resilience increased $\mathbb{R}' > \mathbb{R}$ iff
  - $t_r' \leqslant t_r$ and $t_R' \leqslant t_R$ [work beginning with Cao and Yan]

Complex Network Topology

Multilevel Sprint L1–3 Topology

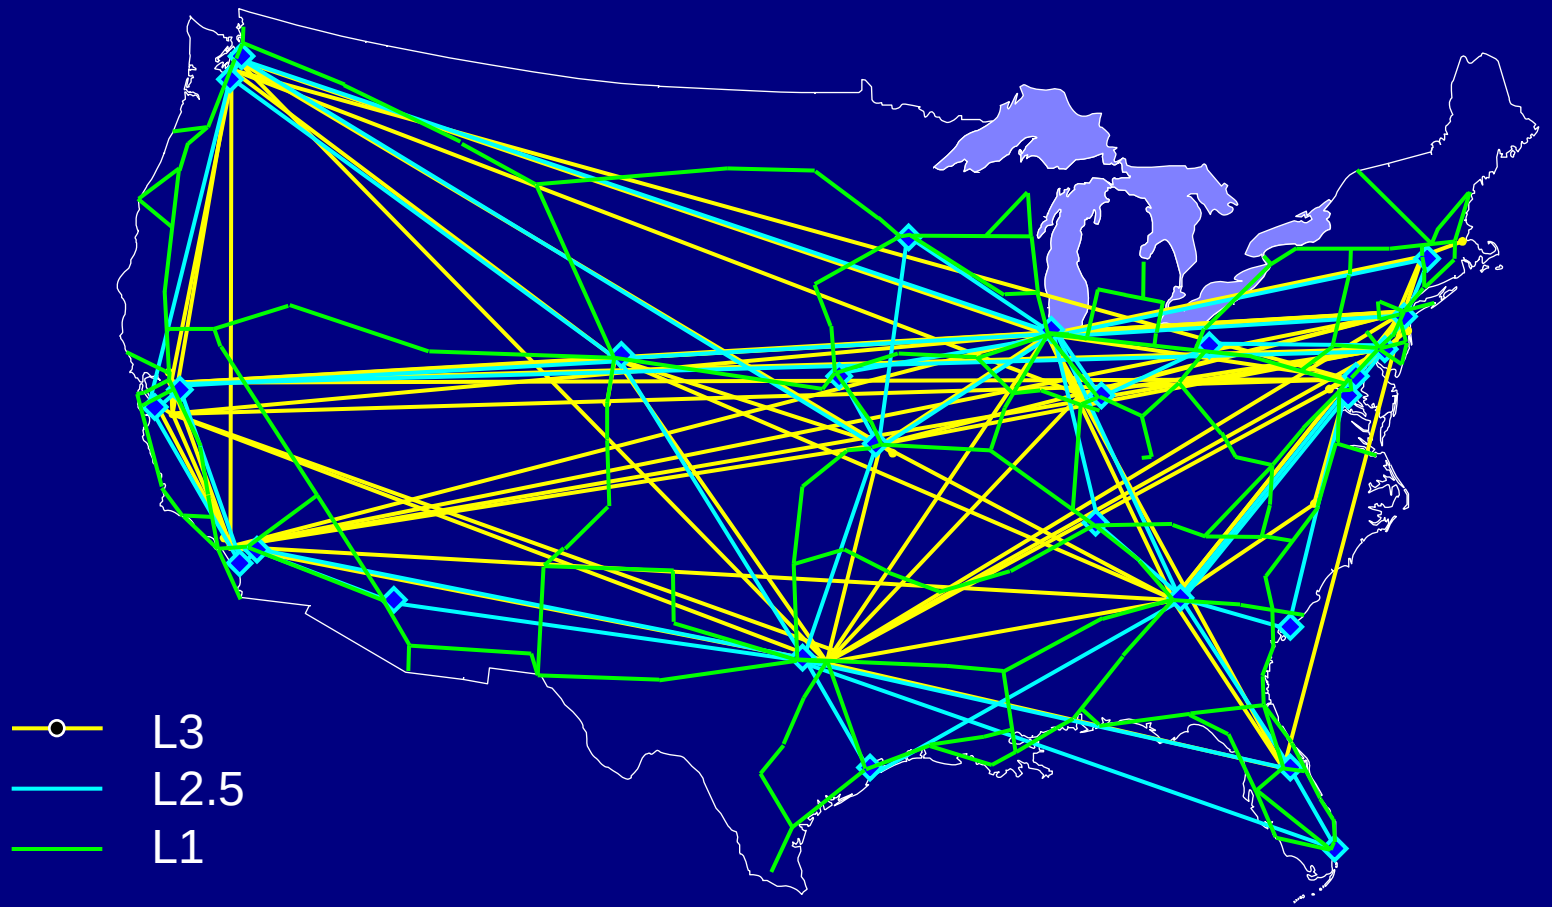

Complex Network Topology

Fiber Relation to Potential Paths

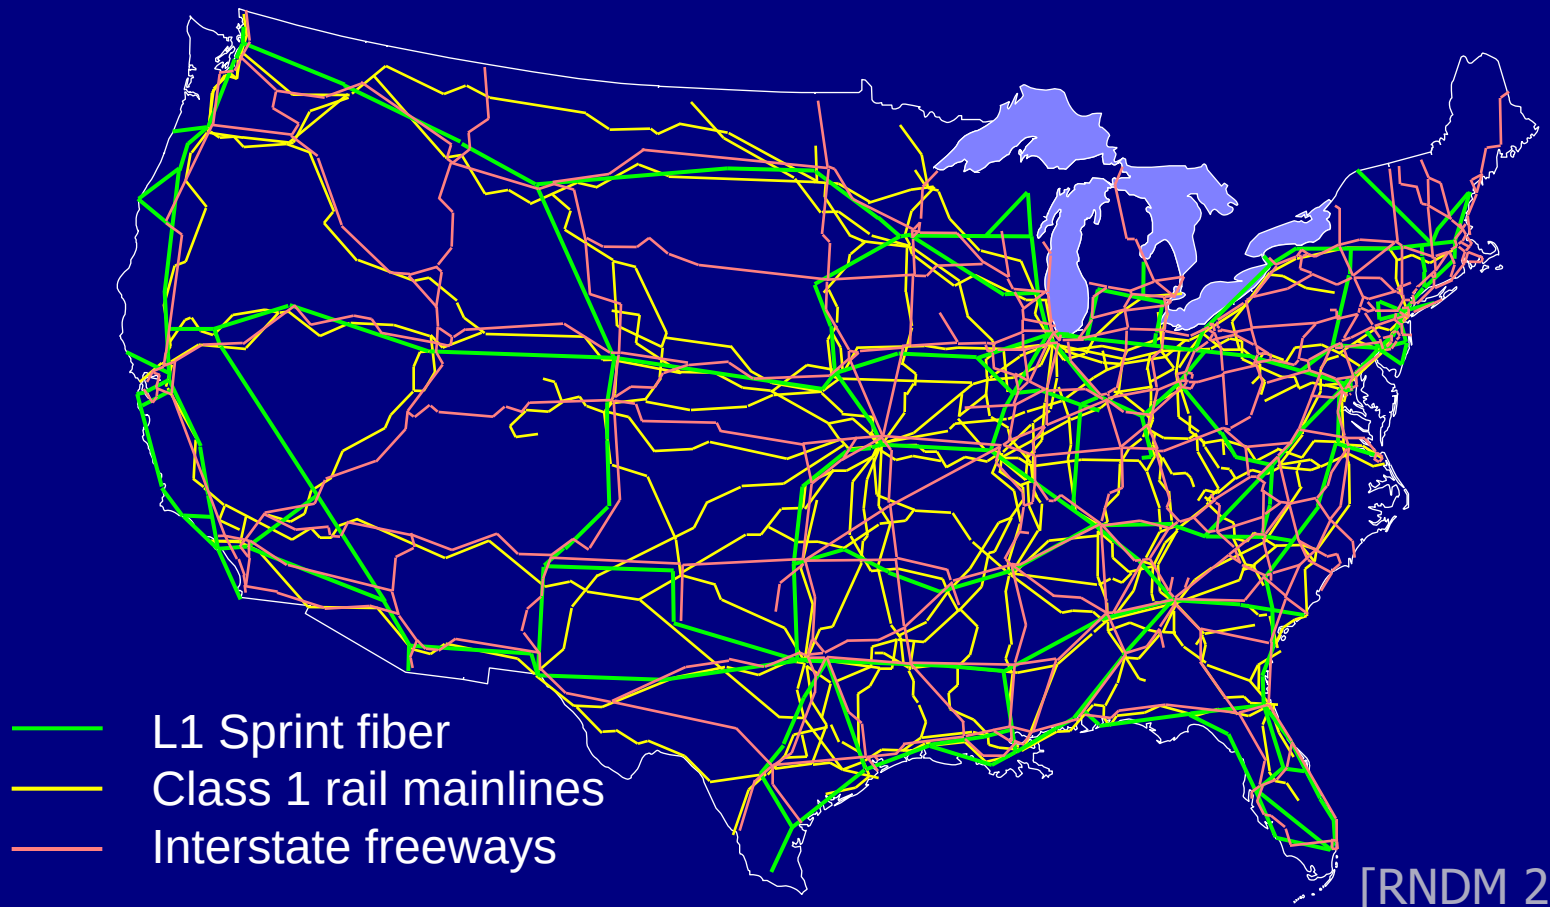

[RNDM 2012]

Network Topology
KU-TopView Topology Viewer

L1 fiber

visualisation ⇔ adjacency matrices

Resilience Analysis
Path and Graph Diversity

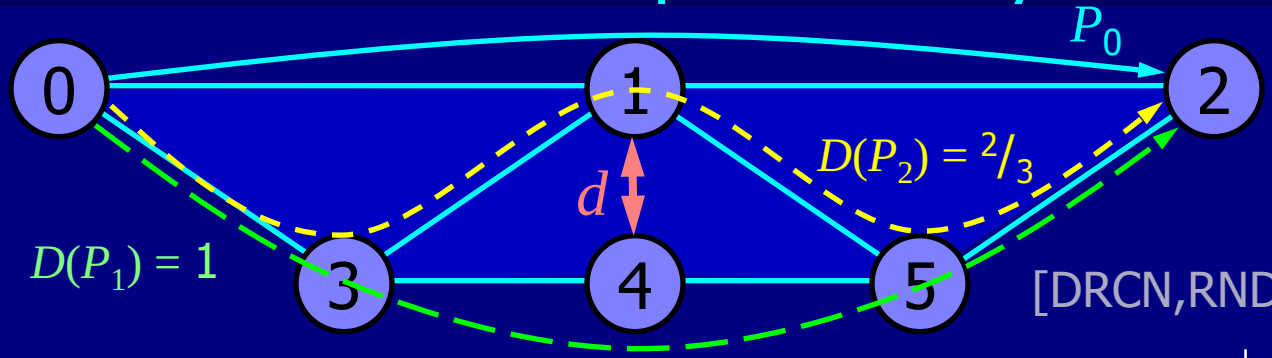

[DRCN,RNDM 2009]

- EPD and TGD
  - diversity in path and graph structure

$$D(P_k) = 1 - \frac{|P_k \cap P_0|}{|P_0|}$$

- cTGD: $e^{\mathrm{TGD}-1} \times h^{-\alpha}$ predictor of resilience

[RNDM 2011, *TSJ* 2012]

- Geographic path diversity
  - distance d between paths beyond source and destination
    - diversity measurement or specification of (k,d)

[NSF with Medhi]

Evaluation Methodology
Challenge Simulation Module

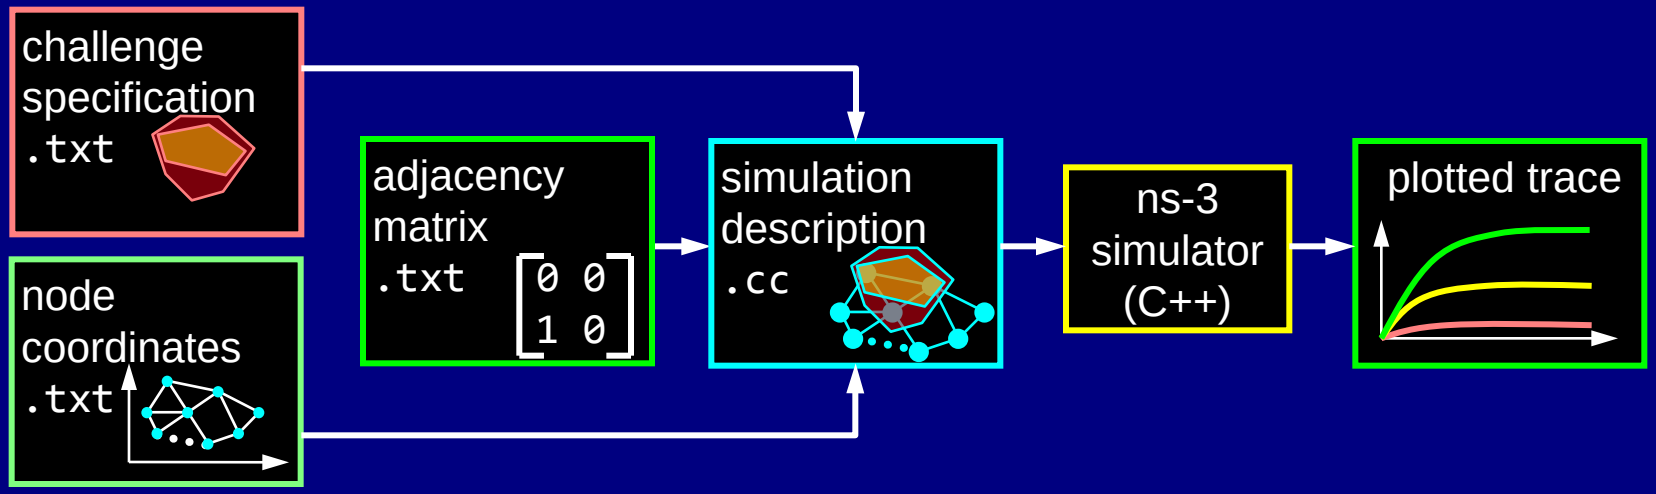

- KU-CSM Challenge Simulation Module
  - challenge specification describes challenge scenario
  - network coordinates provide node geo-locations
  - adjacency matrix specifies link connectivity
  - input to conventional ns-3 simulation run
  - generates trace to plot results with KU-gpWrapper

KU-LoCGen

[RNDM 2010, *TSJ* 2011]

Challenge Modelling
Challenge Types

- Challenge types
  - node or link down
    - random or attack (deg, betweenness, …)
  - area based challenge (disaster)
    - n-sided polygon: (x_0,y_0), … (x_{n-1},y_{n-1})
    - circle centered at (x_0,y_0) with radius r
  - wireless link attenuation or jamming
  - traffic attacks (DoS and DDoS)
- Challenge characteristics
  - type (e.g. wired/wireless)
  - class (e.g. important peering node)
  - dynamic: interval (t_i, t_j), trajectory

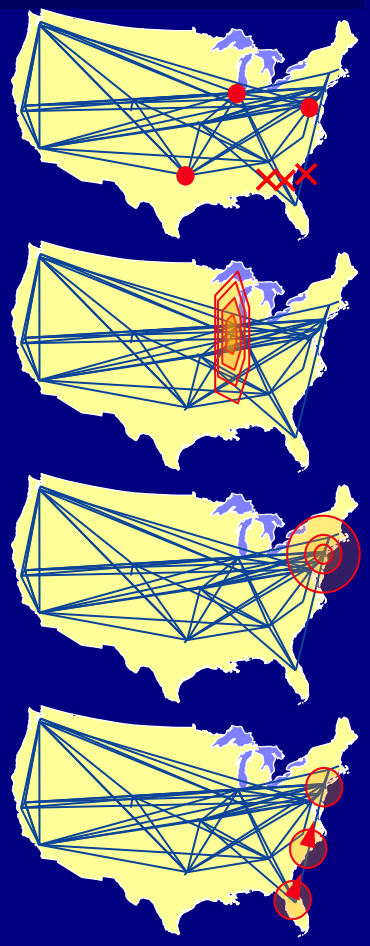

Resilience Evaluation

KU-CSM Physical Disaster Simulation

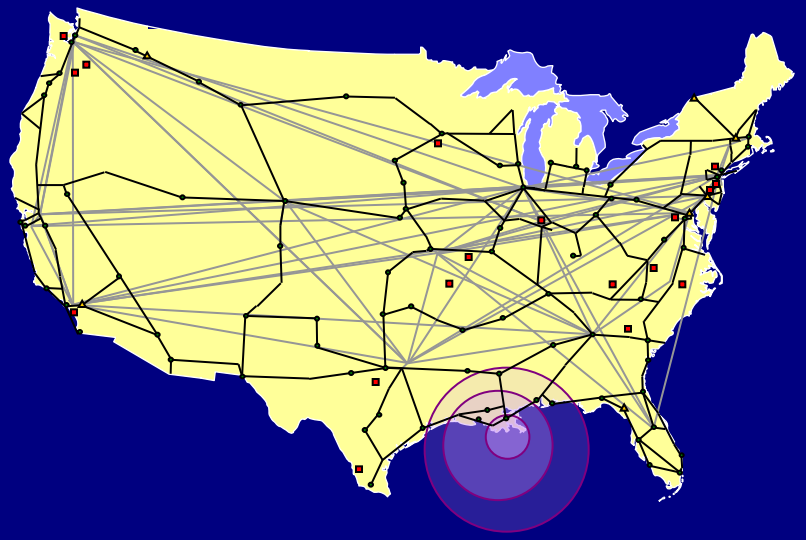

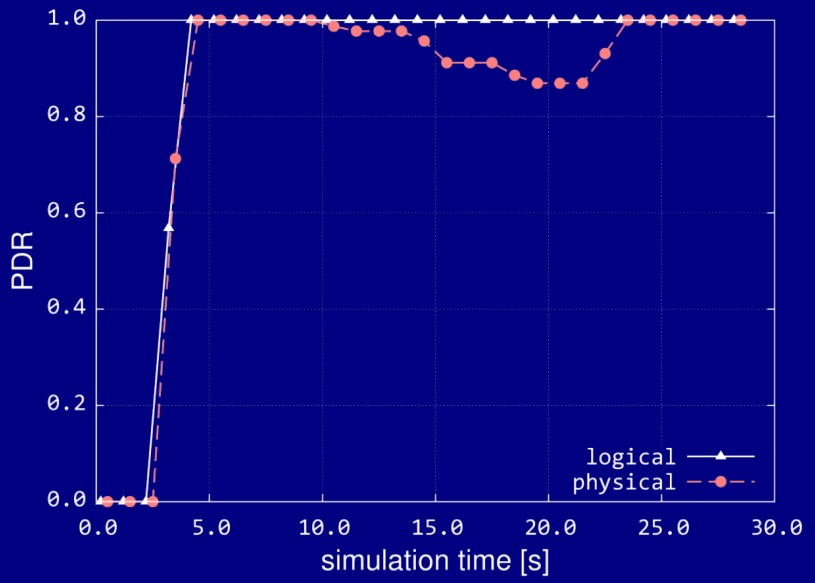

- Example: evolving area-based challenge example
  - circle of increasing size over NOLA (e.g. hurricane)
- Impacts physical infrastructure
  - multilevel analysis measures impact on higher layer services

End

Infrastructure RoW

US Fiber Links

[KMI 1999]

Millimeter-Wave Mesh Networks
Architecture

- Mesh architecture
  - high degree of connectivity
  - alternate diverse paths
    - severely attenuated mm wave
    - alternate mm links
    - alternate lower-freq. RF
    - fiber bypass (competitor)
- Proposed solution
  - route around failures
    - *before* they occur
  - avoid high error links

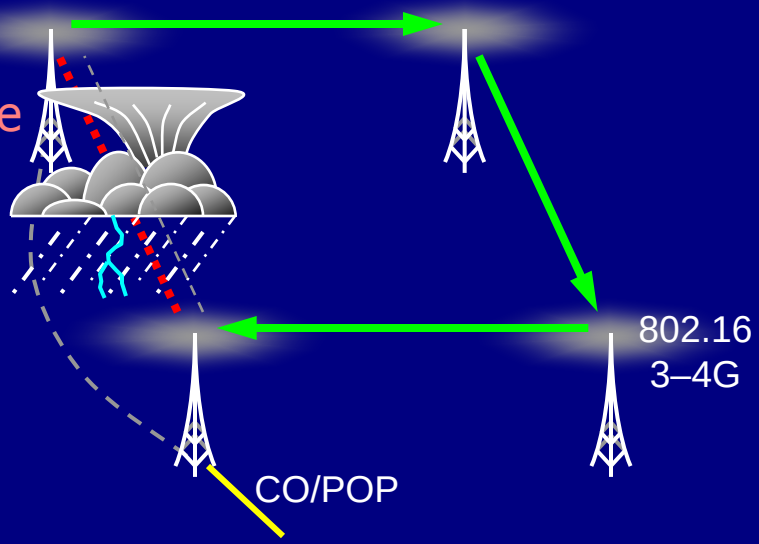

[IWSOS 2008, INFOCOM2009]

Simulations
Observed Storm in Northeast Kansas

- Millimeter-wave grid location
  - 38.8621N, 95.3793W
- Storm observed at:
  - 20:39:26Z 30 Sep 2008

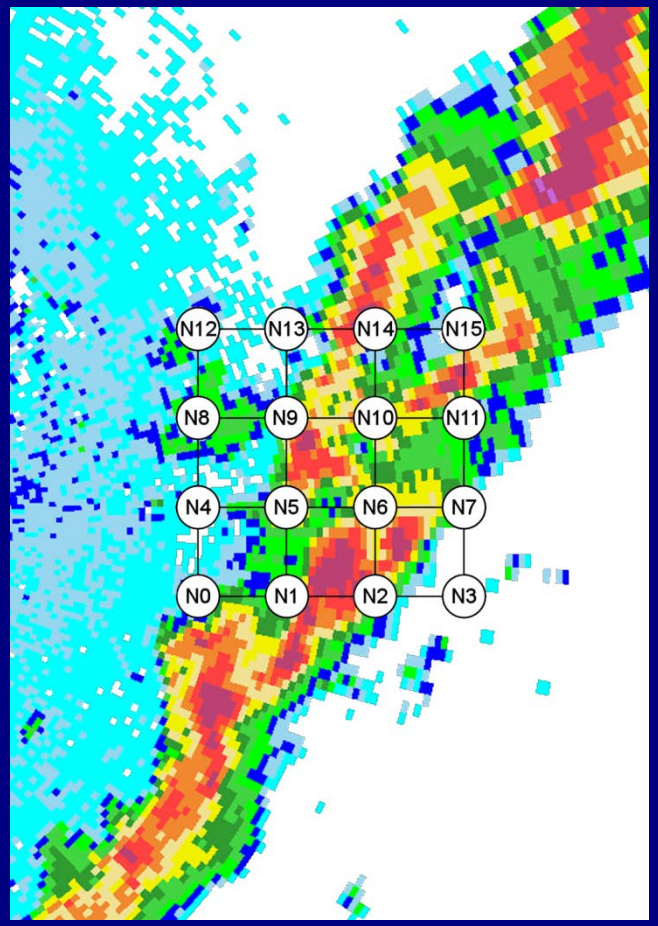

Synthetic Storm

Performance Analysis: Packet Loss