Geodiverse Multipath Communication
with Structural Multilevel Diversity for the Resilient Future Internet

James P.G. Sterbenz*†
Джеймс Ф.Г. Штербэнз　제임스 스털벤츠　司徒傑莫
Dongsheng Zhang 张东升, Yufei Cheng 成宇飞,
Anh Nguyễn, Mohammed Alenazi
*Department of Electrical Engineering & Computer Science
Information Technology & Telecommunications Research Center
The University of Kansas
†School of Computing and Communications, Infolab 21
Lancaster University

jpgs@ittc.ku.edu

http://www.ittc.ku.edu/~jpgs
http://wiki.ittc.ku.edu/resilinets

© 2014 James P.G. Sterbenz

Where is Kansas?

Geography Lesson

Multilevel Structural Diversity
Outline

- ResiliNets review
- Multilevel interrealm resilience
  - resilience to attackers
  - resilience to large scale disasters

Resilience and Survivability

Motivation and Definition

- Increasing reliance on network infrastructure
  - ⇒ Increasingly severe consequences of disruption
  - ⇒ Increasing attractiveness as target from bad guys
- Need *resilience*
  - provide and maintain acceptable service
  - in the face of faults and challenges to normal operation
- Challenges
  - …
  - large scale disasters (natural and human-caused)
  - malicious attacks from intelligent adversaries

ResiliNets Initiative
Goals

- Understand network structure and vulnerabilities
  - develop new models and tools for analysis
- Develop ways to increase network resilience
  - improving existing networks under cost constraints
  - increase cost to attackers
  - Future Internet design
  - validate by analysis, simulation, and experimentation
- Funded primarily by
  - US NSF FIND and GENI programs and open call (with Medhi)
  - US DoD
  - EU FP6 and FP7 FIRE programme (with David Hutchison)

ResiliNets Strategy

D^2R^2 + DR

- Two phase strategy for resilience
- Real time control loop: D^2R^2
  - defend
    - passive
    - active
  - detect
  - remediate
  - recover
- Background loop: DR
  - diagnose
  - refine

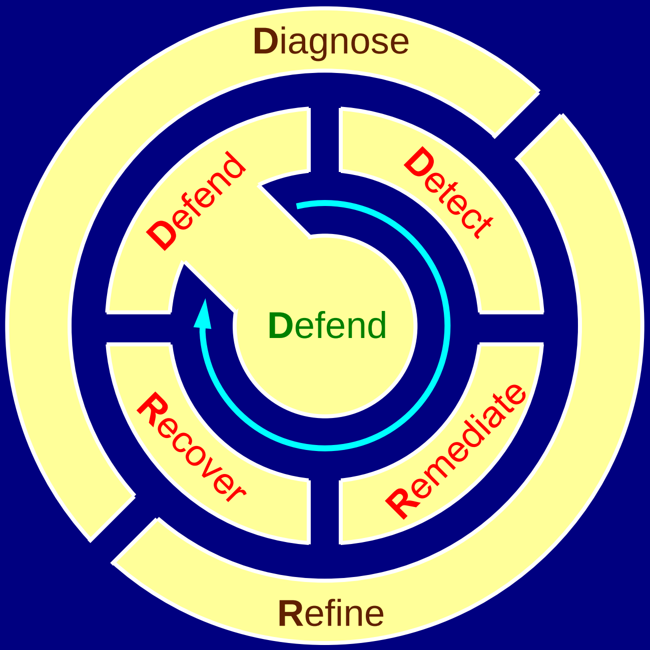

[Wiki 2005, *ComNet* 2010]

ResiliNets Principles
High Level Grouping

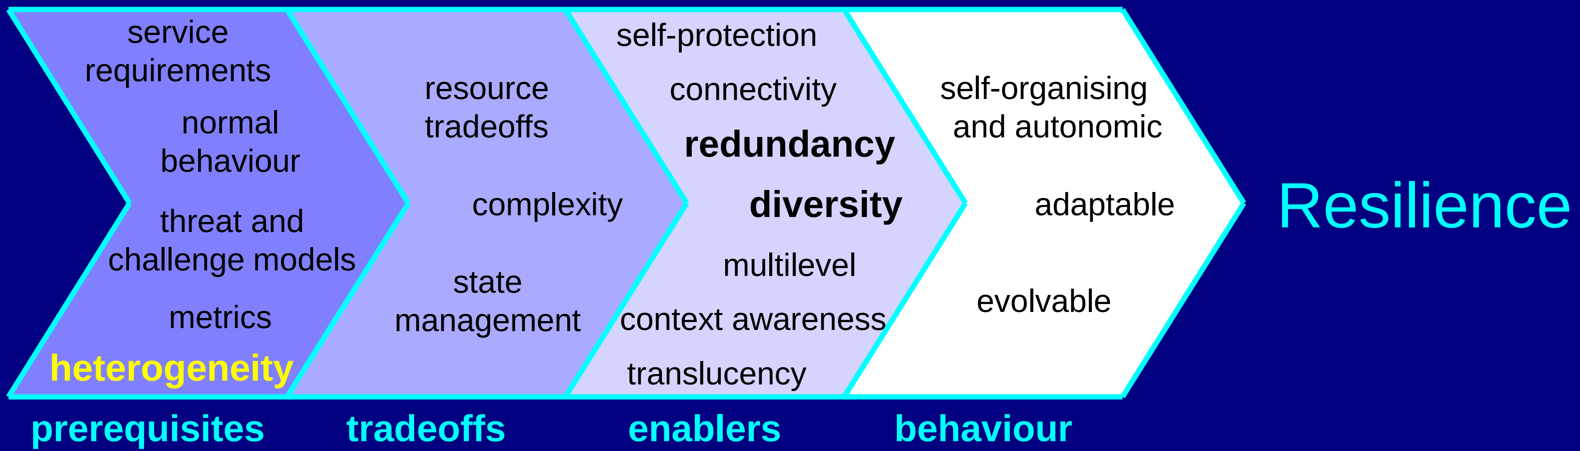

- Prerequisites: to understand and define resilience
- Tradeoffs: recognise and organise complexity
- Enablers: architecture and mechanisms for resilience
- Behaviour: require significant complexity to operate

Resilience Principles
Redundancy, Diversity, Heterogeneity

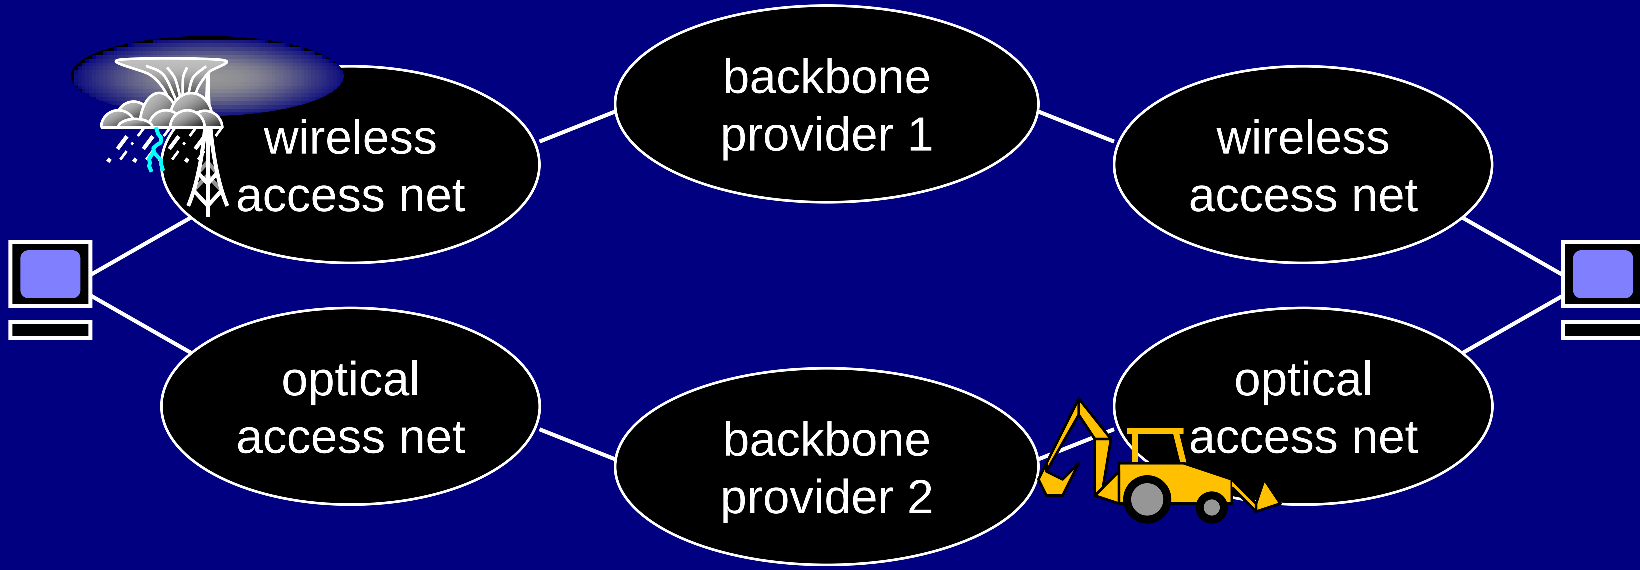

- Diversity
  - mechanism (wired & wireless), provider, *geographic path*
- Multipath transport
  - spreading (erasure code) or as hot-standby

Multilevel Structural Diversity
Multilevel Interrealm Resilience

- ResiliNets review
- Multilevel interrealm resilience
  - resilience to attackers
  - resilience to large scale disasters

Multilevel Network Topology

Example: Sprint L3 IP PoP Topology

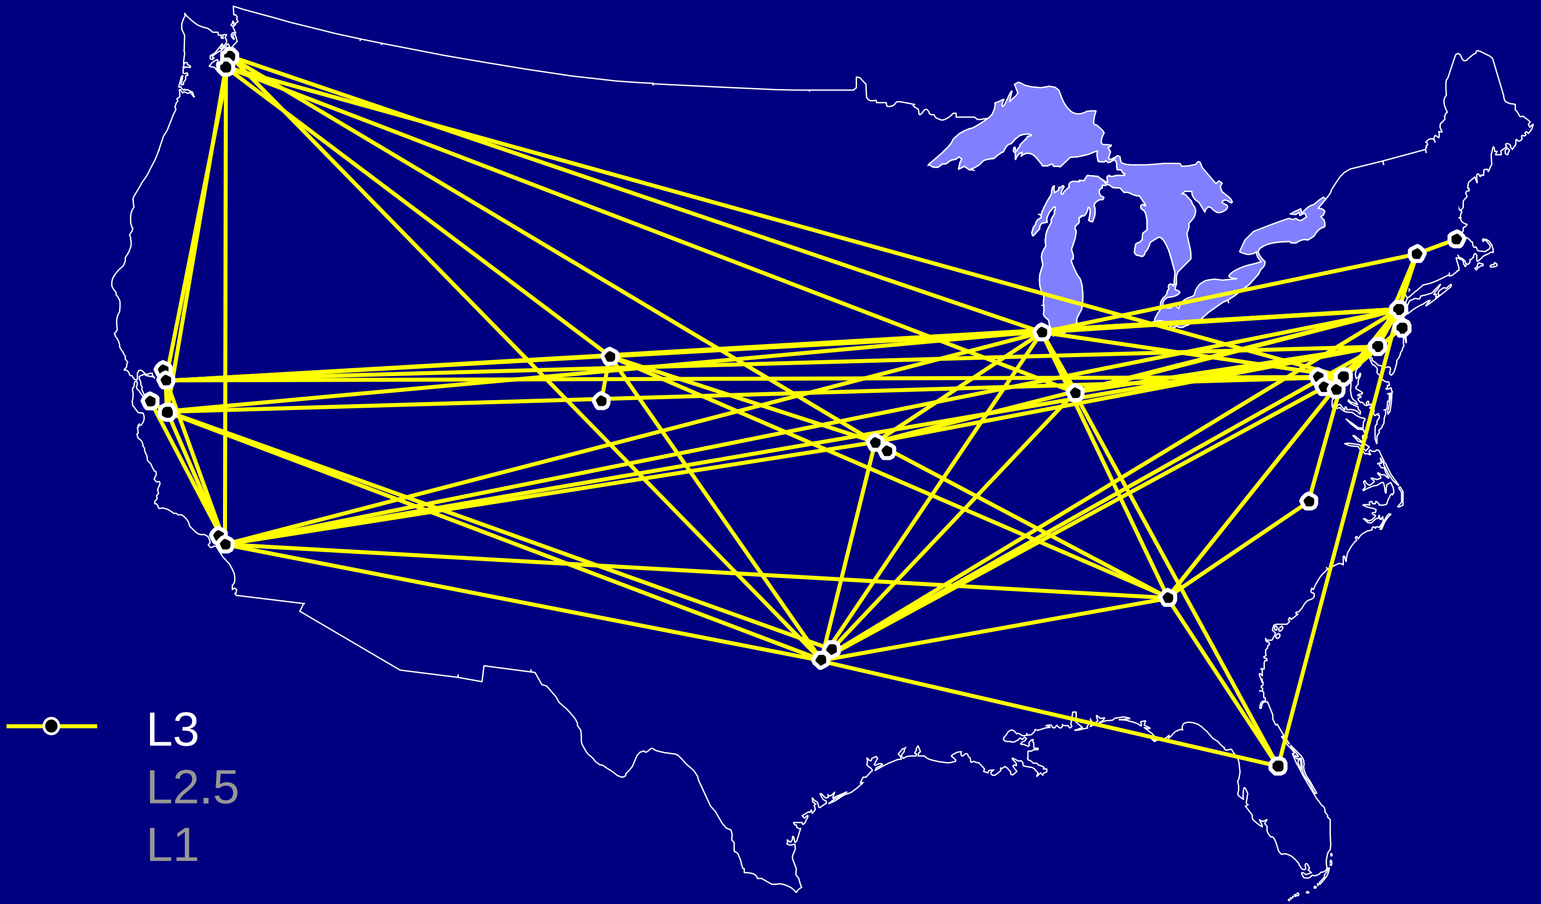

Multilevel Network Topology

Example: Sprint L3 overlay on L2.5

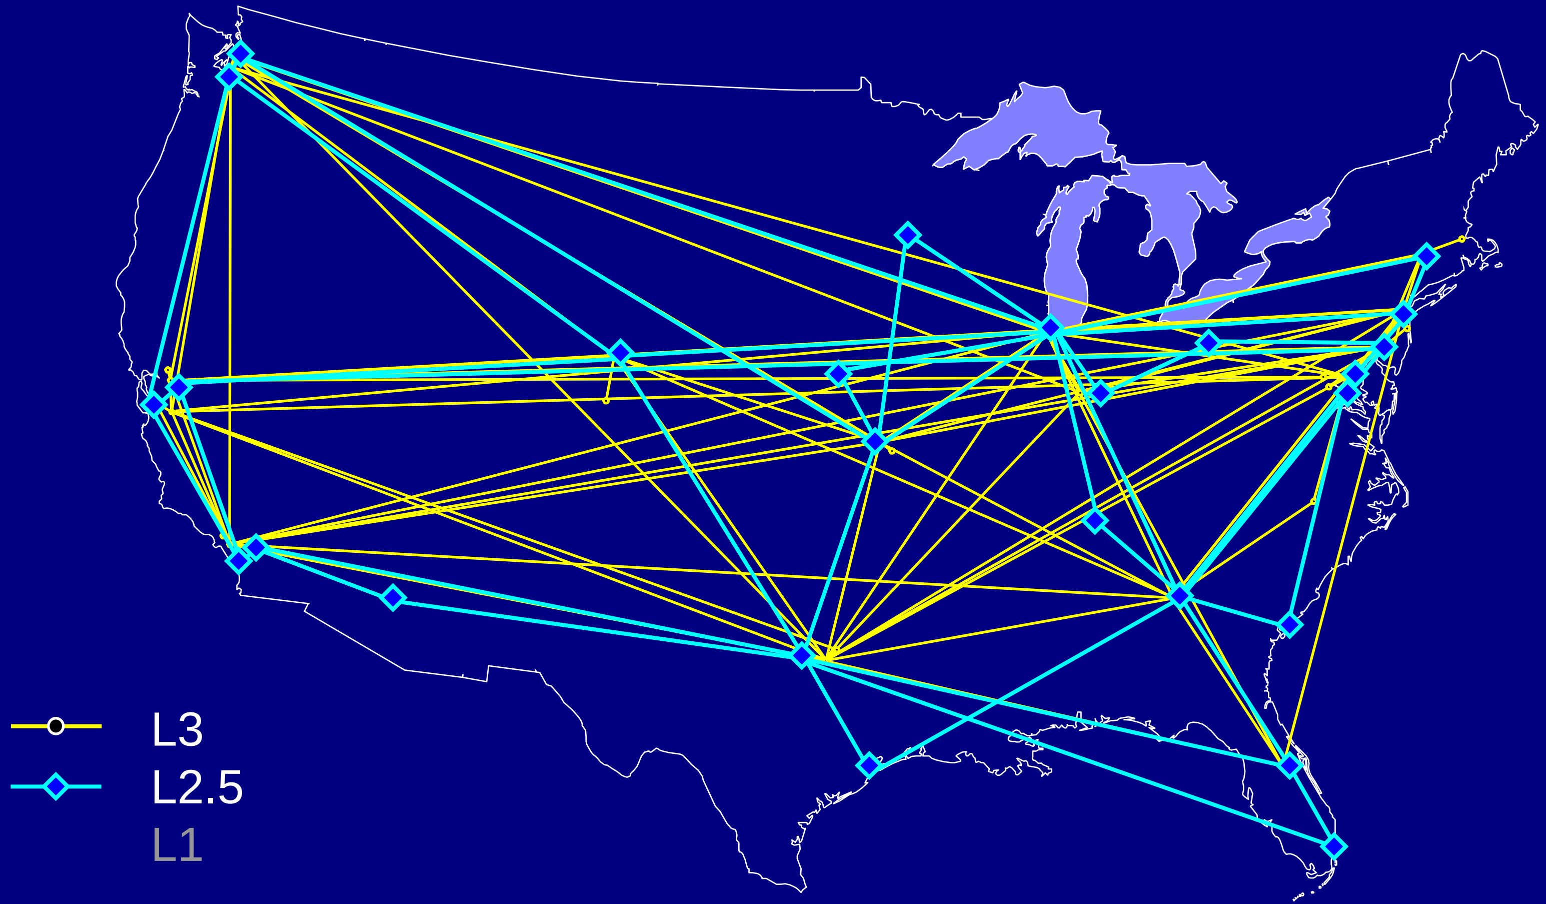

Multilevel Network Topology

Example: Sprint L2.5 MPLS PoP Topology

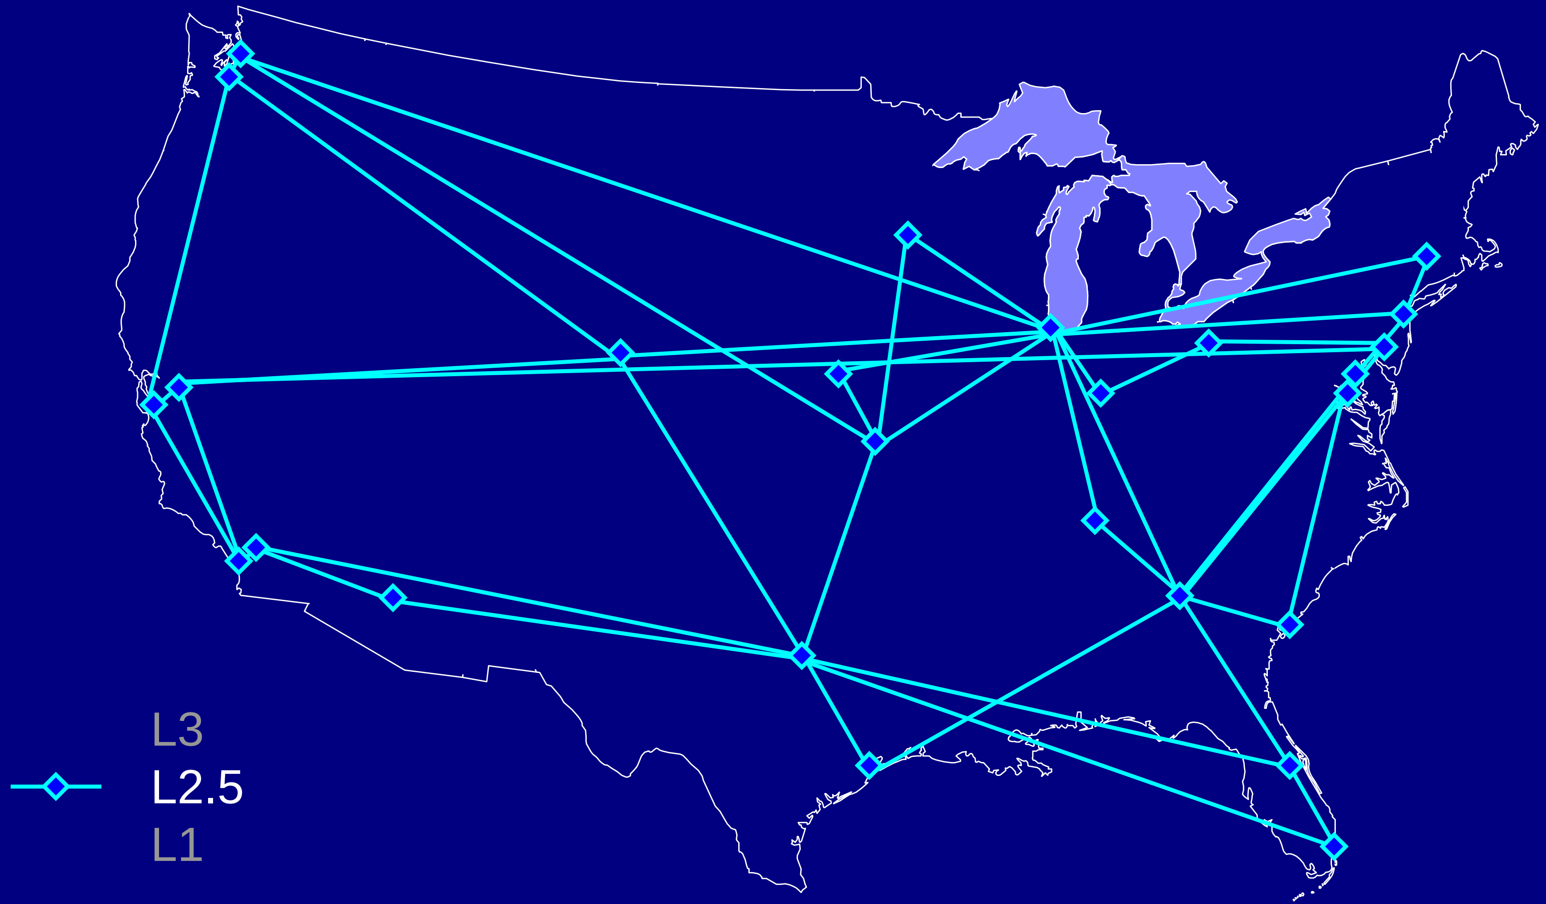

Multilevel Network Topology

Example: Sprint L2.5 overlay on L2/1

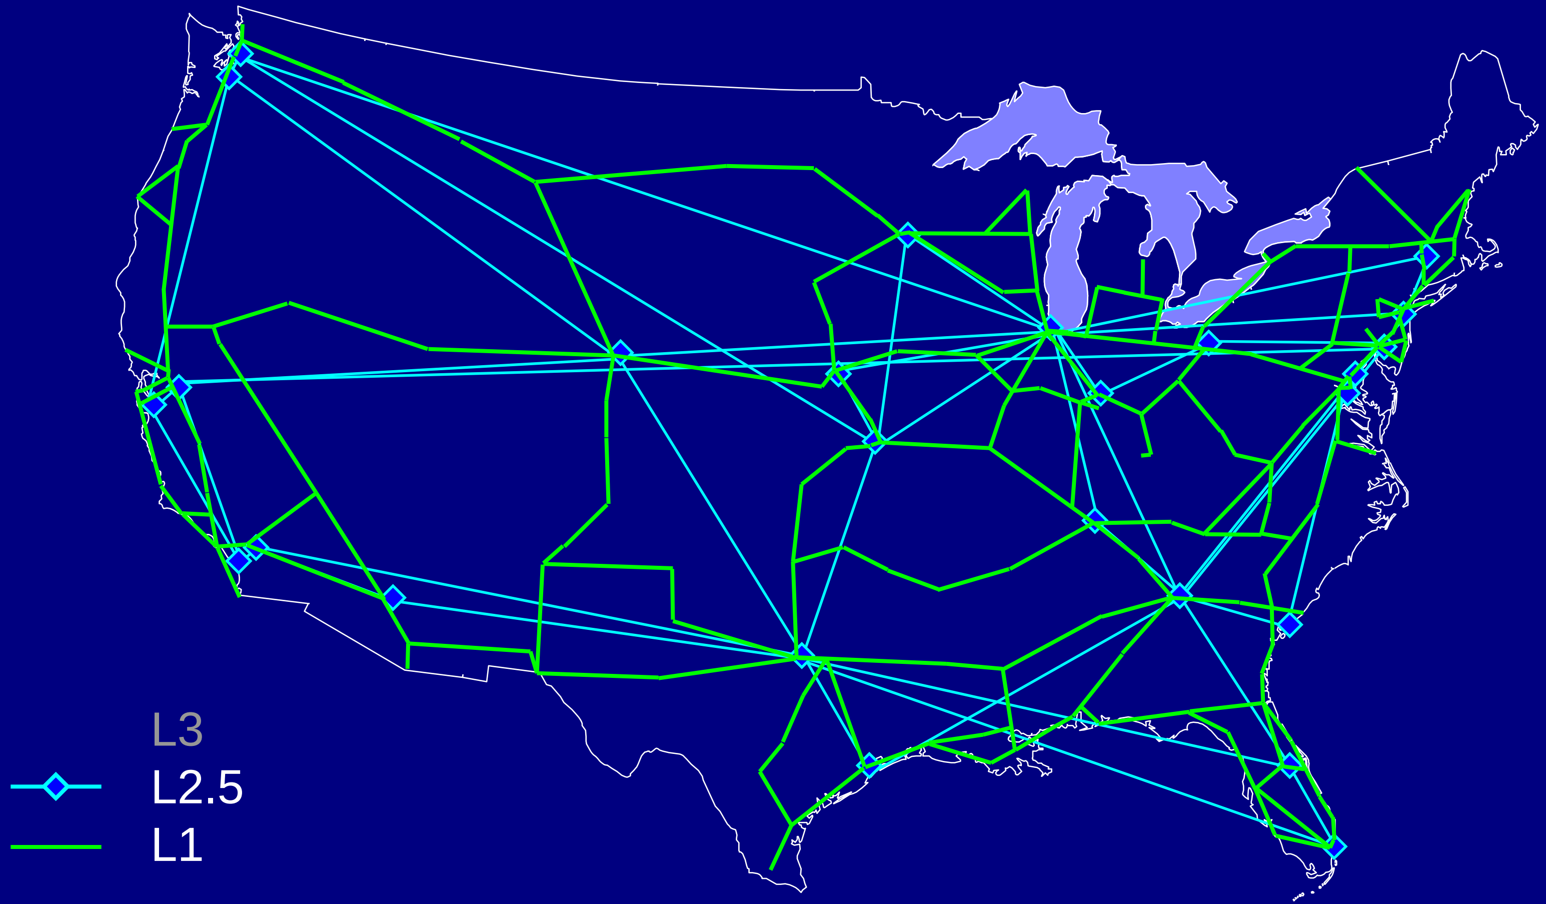

Multilevel Network Topology

Example: Sprint L1 Physical Fiber Topology

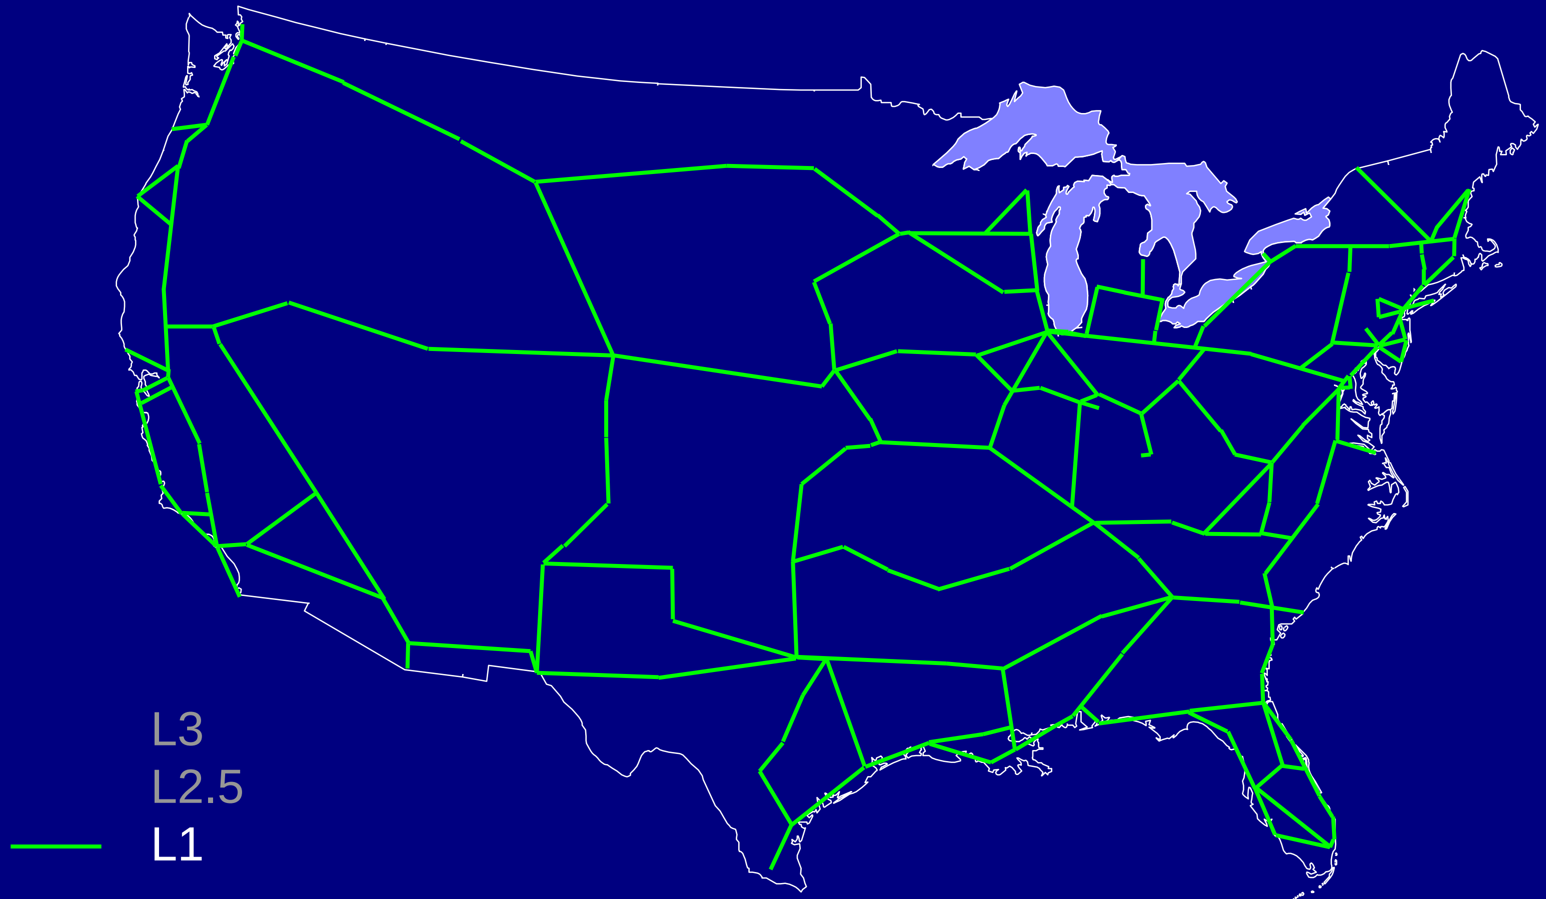

Multilevel Network Topology

Example: Sprint L1–3 Topology

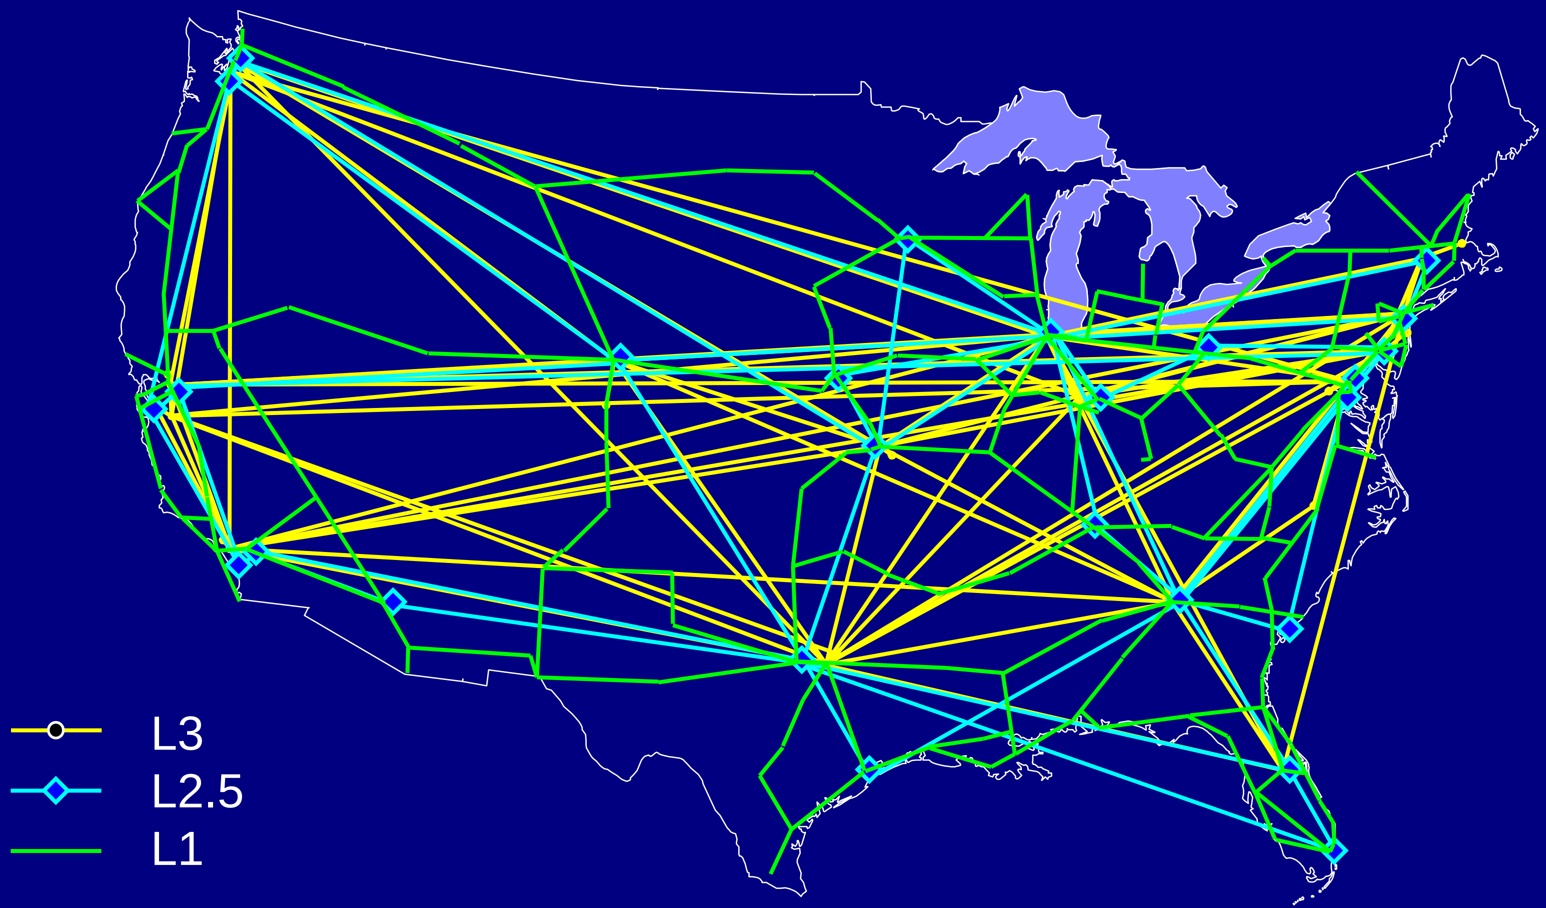

Complex Network Topology

KU-TopView Topology Viewer

L1 Sprint fiber

visualisation ⇔ adjacency matrices

Challenge Simulation

KU-CSM

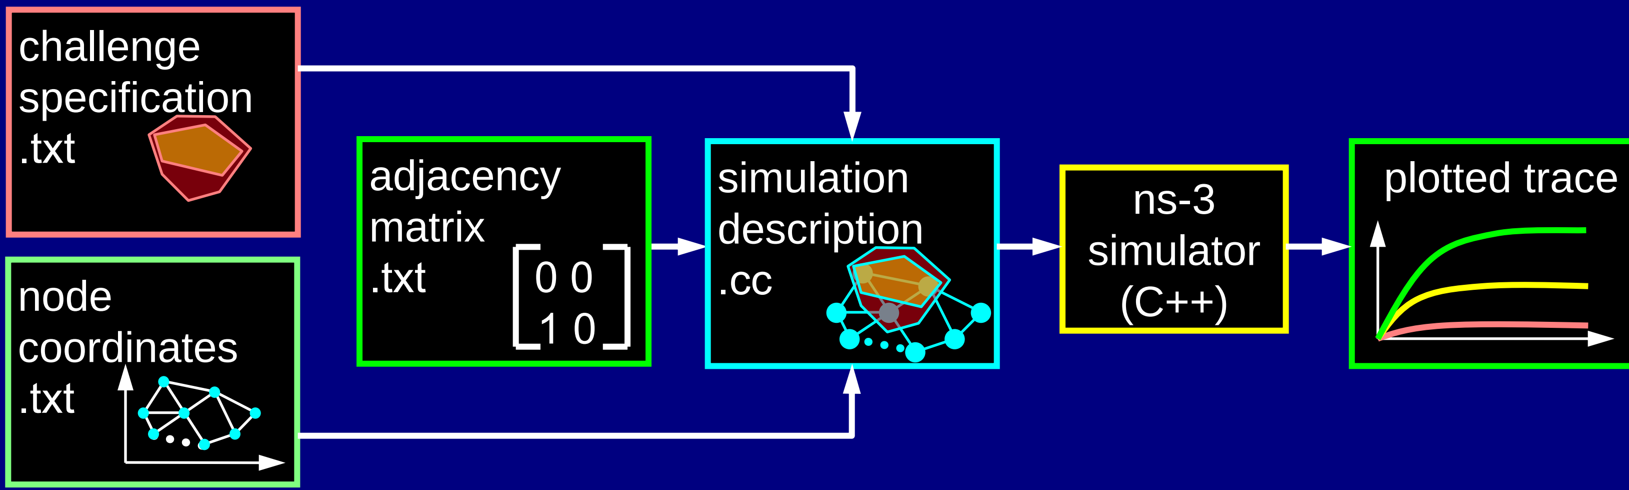

- KU-CSM Challenge Simulation Module
  - challenge specification describes challenge scenario
  - network coordinates provide node geo-locations
  - adjacency matrix specifies link connectivity KU-LoCGen
  - input to conventional ns-3 simulation run
  - generates trace to plot results with KU-gpWrapper [RNDM 2010]

Challenge Simulation
Challenge Types

- Challenge types
  - node or link down
    - random or attack (deg, betweenness, …)
  - area based challenge
    - 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

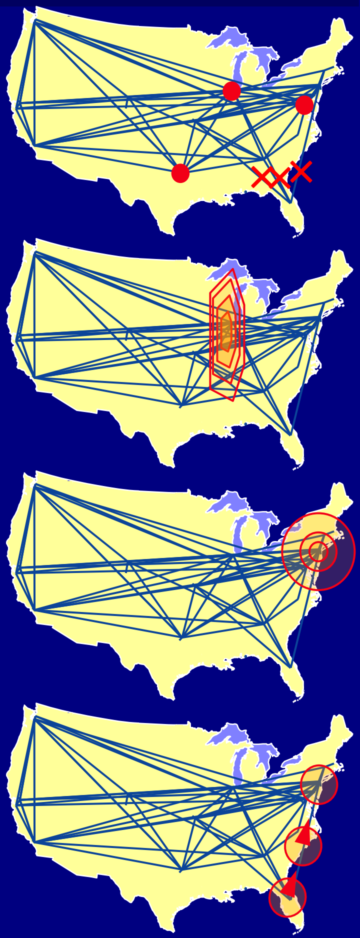

Multilevel Structural Diversity
Resistance to Attackers

- ResiliNets review
- Multilevel interrealm resilience
  - resilience to attackers
  - resilience to large scale disasters

Multilevel Network Analysis
Abstraction of Internet Topology

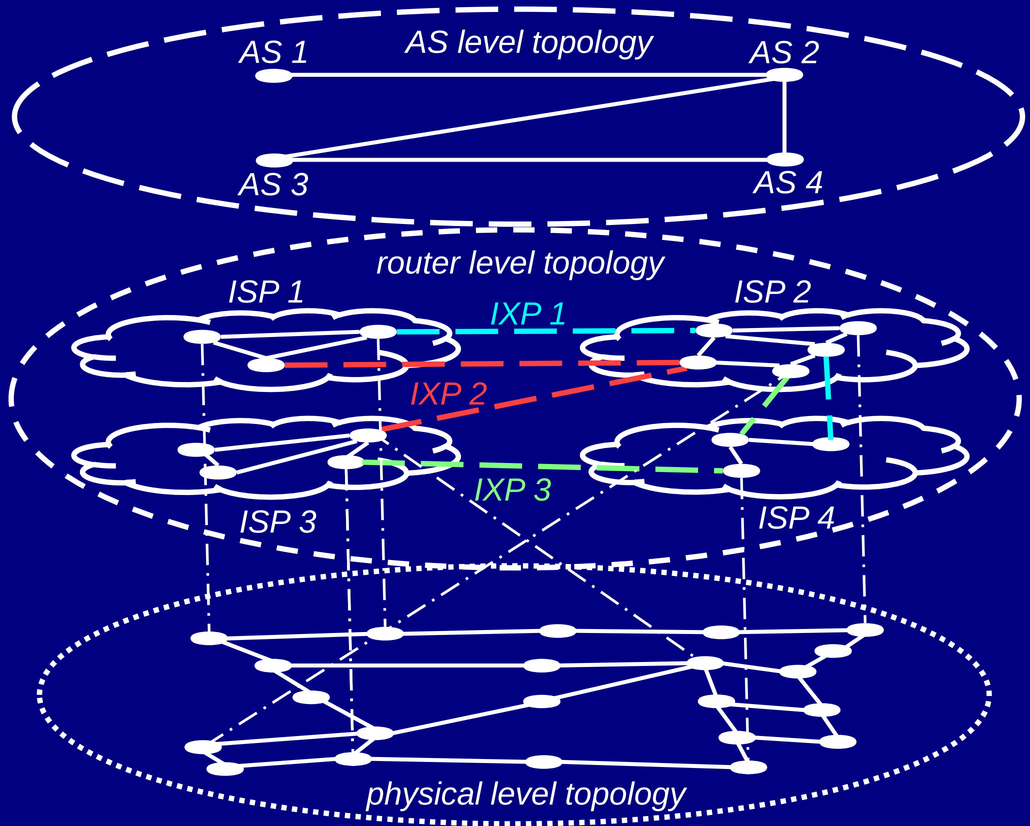

[DRCN 2013]

Multilevel Network Analysis

Multilevel Graph Model

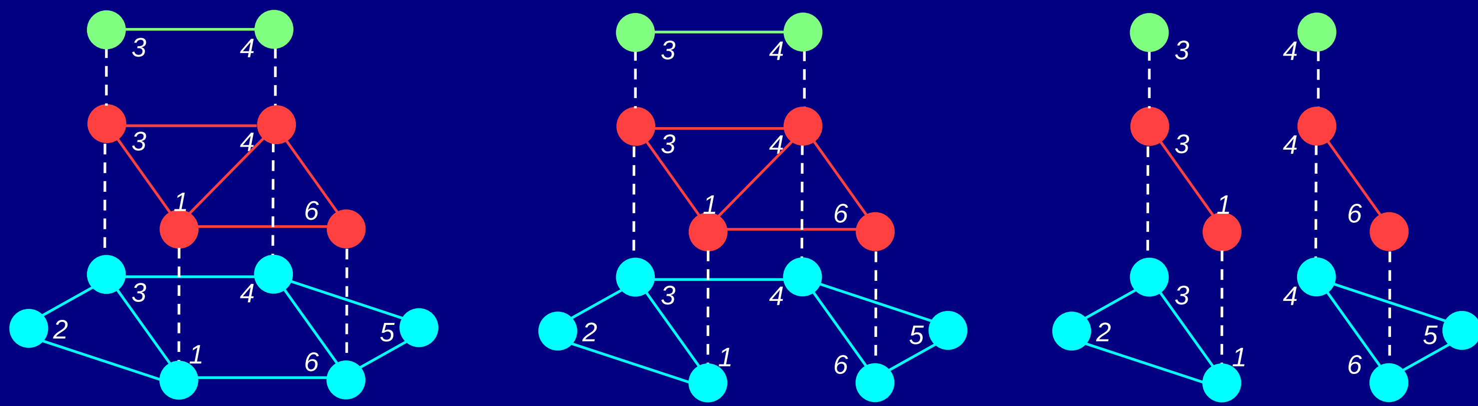

Connected network *Disconnected network* *Partitioned network*

- Multilevel model for unweighted & undirected graphs
- Two requirements for multilevel graph model:
  - nodes at the above level are subset of lower level
  - nodes that are disconnected below are disconnected above

Resilience Analysis

Graph-Theoretic Properties of Networks

Topology	Sprint Physical	Sprint Logical	AT&T Physical	AT&T Logical	US Highways
Number of nodes	263	28	361	107	400
Number of links	311	76	466	140	540
Maximum degree	6	14	7	23	7
Average degree	2.37	5.43	2.58	2.62	2.7
Degree assortativity	-0.17	-0.23	-0.16	-0.4	0.11
Node closeness	0.07	0.48	0.08	0.3	0.08
Clustering coefficient	0.03	0.41	0.05	0.09	0.05
Algebraic connectivity	0.0053	0.6844	0.0061	0.1324	0.0059
Network diameter	37	4	37	6	40
Network radius	19	2	19	3	21
Average hop count	14.78	2.19	13.57	3.38	13.34
Node betweenness	11159	100	15970	2168	22798
Link betweenness	9501	27	14270	661	18585

Multilevel Resilience

Effect of Physical Failures on L3 Topology

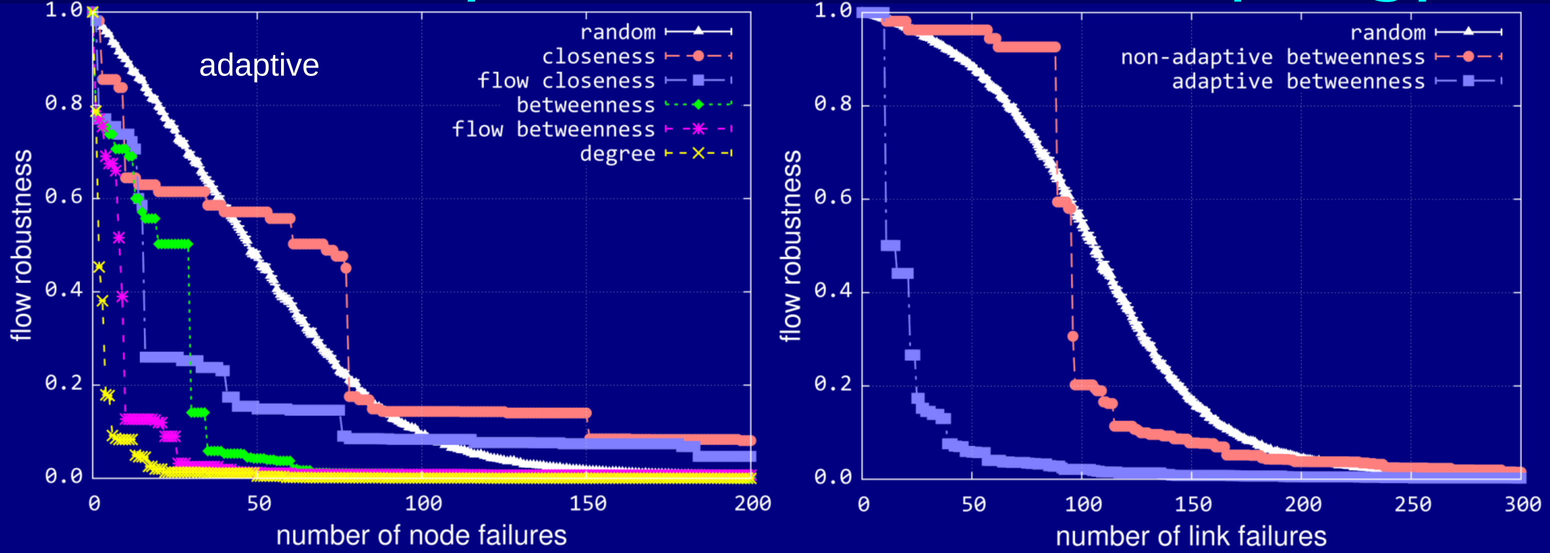

- Attacks against physical infrastructure
  - based on centrality (importance) metrics
  - adaptive recomputes metrics after each node failure)
- Analysis of impact on higher layer flows
  - heuristics to add elements under cost constraints

Multilevel Structural Diversity
Resilience to Large-Scale Disasters

- ResiliNets review
- Challenge Taxonomy
- Multilevel interrealm resilience
  - resilience to attackers
  - resilience to large scale disasters

Simulation Analysis

Example: Multilevel Analysis of Disaster

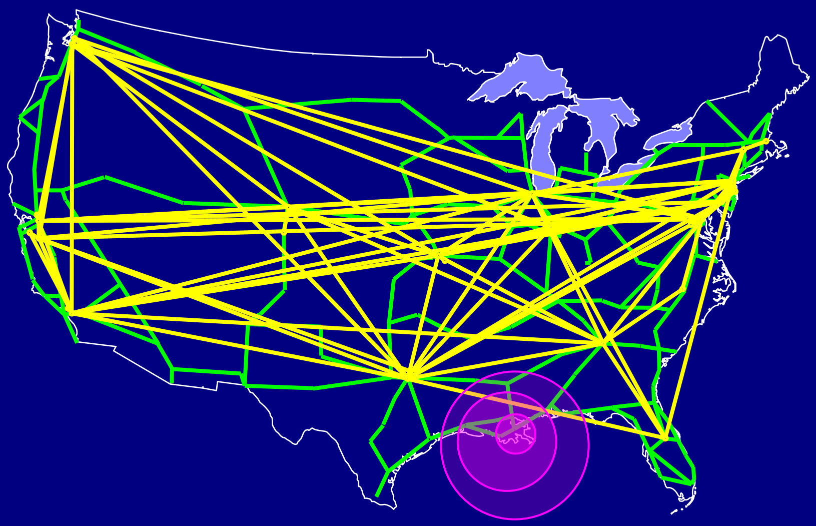

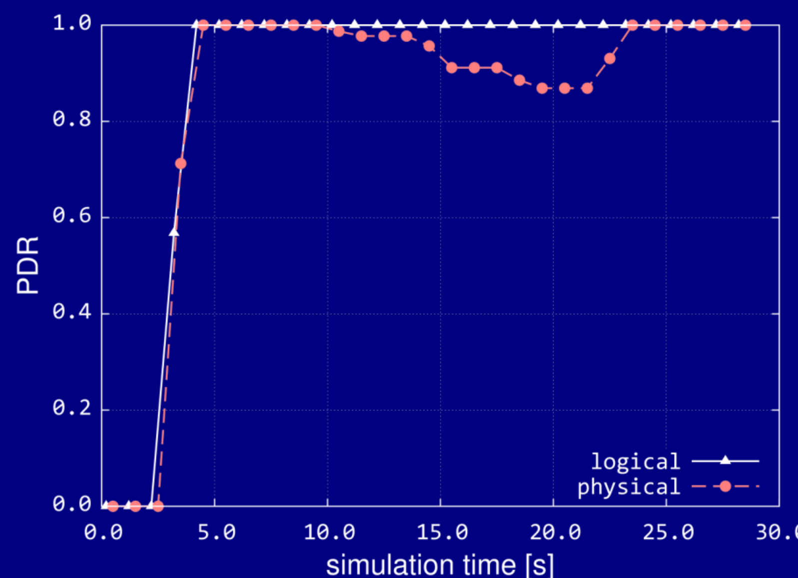

- Hurricane disaster in New Orleans area
- Destruction of physical infrastructure
- Effect on IP-layer network services

Resilience Analysis
Path and Graph Diversity

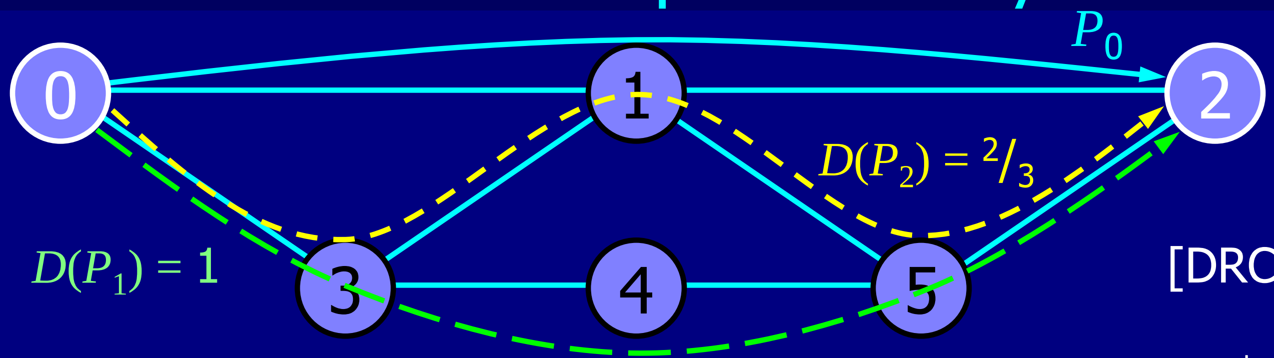

[DRCN 2009]

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

- Path diversity
  - measure of links and nodes in common
- EPD: effective path diversity [0,1)
  - normalised diversity with respect to a single shortest path
  - measure of E2E flow resilience
- TGD: total graph diversity is average of EPD
  - for all pairs: quantifies available diversity in graph

Resilience Analysis
Path and Graph Diversity with Distance Metric

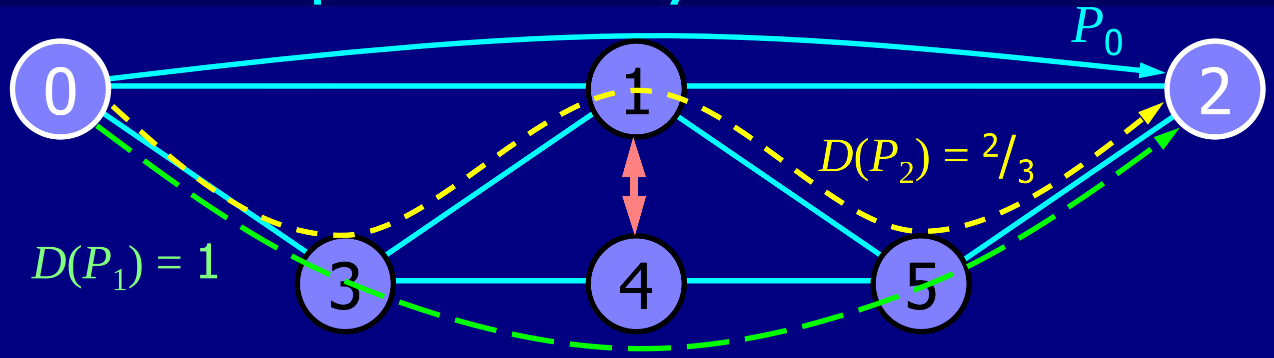

- cTGD: compensated TGD
  - weighted to be predictive of flow robustness [RNDM 2010]
  - algebraic connectivity also fair predictor or flow robustness
- GeoPath diversity
  - distance d between paths beyond source and destination
  - GeoResLSR: $(k, d, [s,t])$ multipath geographic routing
    - number of paths k

Resilience Analysis

Compensated Total Graph Diversity

Metric Network	surv	deg	c TGD	TGD	clus coef	dia	hop cnt	clse	nod btw	link btw
full mesh	01	01	01	01	01	01	01	01	01	01
Level3	02	02	02	02	02	04	02	03	10	09
AboveNet	03	03	03	08	03	03	03	02	05	03
. . .										
ring	15	07	15	13	15	09	15	15	04	08
AT&T L1	16	07	16	03	13	10	16	16	15	17
Sprint L1	17	07	17	06	14	10	17	17	12	16

- cTGD much better predictor of flow robustness
  - cTGD with $\alpha = 0.25$ perfect predictor for these 17
    - 13 real networks plus 4 regular topologies

[RNDM 2011]

ResiliNets Protocols
Cross-Layer Model: Generic

- *Knobs* $\mathbb{K}_{i \to i-1} = \{k_i\}$ influence behaviour to levels below
- *Dials* $\mathbb{D}_{i+1 \leftarrow i} = \{d_i\}$ expose characteristics to upper levels
- Levels (of significance to ResiliNets)
  - 8: social
  - 7: application
  - 4: end-to-end transport
  - 3i: inter-realm (domain)
  - 3r: routing
  - 3t: logical topology
  - 2: hop-by-hop links
  - 1: physical topology

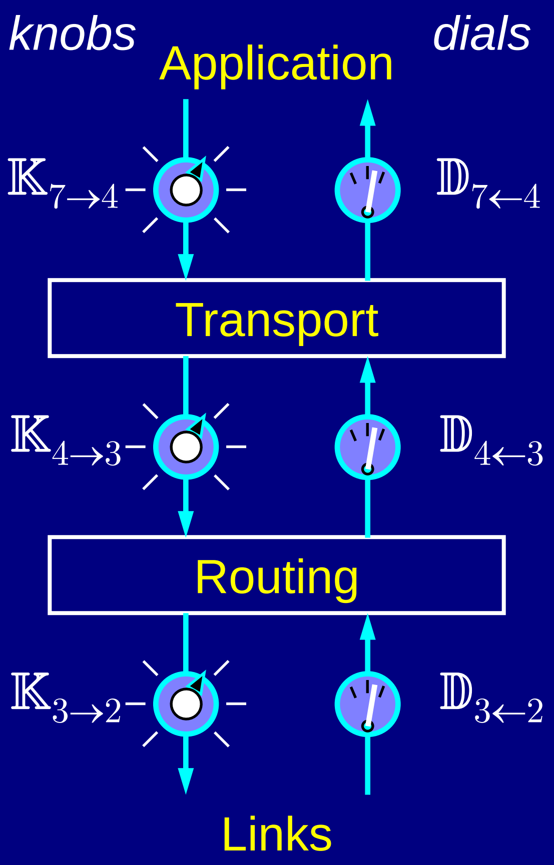

Resilient Transport: ResTP
Overview

- ResTP: Resilient Transport Protocol
  - flexible and composable [ala TP++ [Feldmeier, McAuley]
- Flexible and composable
  - flow setup and management
    - including multipath support
  - error control
  - transmission (flow and congestion) control
- Cross-layered
  - applications specify service and threat model
  - behaviour based on path characteristics
  - specifies path requirements to GeoDivRP

Resilient Transport: ResTP
Reliability Modes

- Reliability (combination of flow and error modes)
  - full reliability: E2E 3-way handshake and ACKs
  - nearly-reliable: custody transfer at GW with e2e ACKs
  - quasi-reliable: E2E FEC giving statistical reliability
  - none (flow): connection oriented best effort
  - none (datagram): connectionless best effort (UDP-like)
- Chosen using cross-layering
  - service specification and threat model from application
  - path characteristics from lower layers

Resilient Transport: ResTP

Flow Modes

- Multiple flow modes
  - hard connections (3-way handshake)
  - opportunistic connections (signalling overlaps data)
  - custody transfer at realm boundaries (for DTNs)
    - AeroTP subset of ResTP uses this
  - soft-state flows
  - signalled flow with datagrams
  - individual datagrams

Resilient Transport: ResTP
Error Control

- Multipath modes
  - alternate path added on-demand
  - alternate path as hot-standby
  - erasure coding across k paths (typically $k=3$)
    - best coding for large skew?
- Per subflow modes
  - ARQ for reliable service
    - SACK, MACK, NAK, SNACK (SCPS-style)
  - HARQ for reliable service on lossy path
  - adaptive FEC for quasireliable service
  - none for unreliable service

Resilient Transport: ResTP
E2E Transport vs. HBH Error Control

- Alternatives
  - N none
  - O open loop (FEC)
  - C closed loop (ARQ)
    - S&W, GB-*N*, SelRep
- Location
  - – HBH
  - – E2E
- App requirements
  - – unreliable
  - – quasi-reliable
  - – reliable

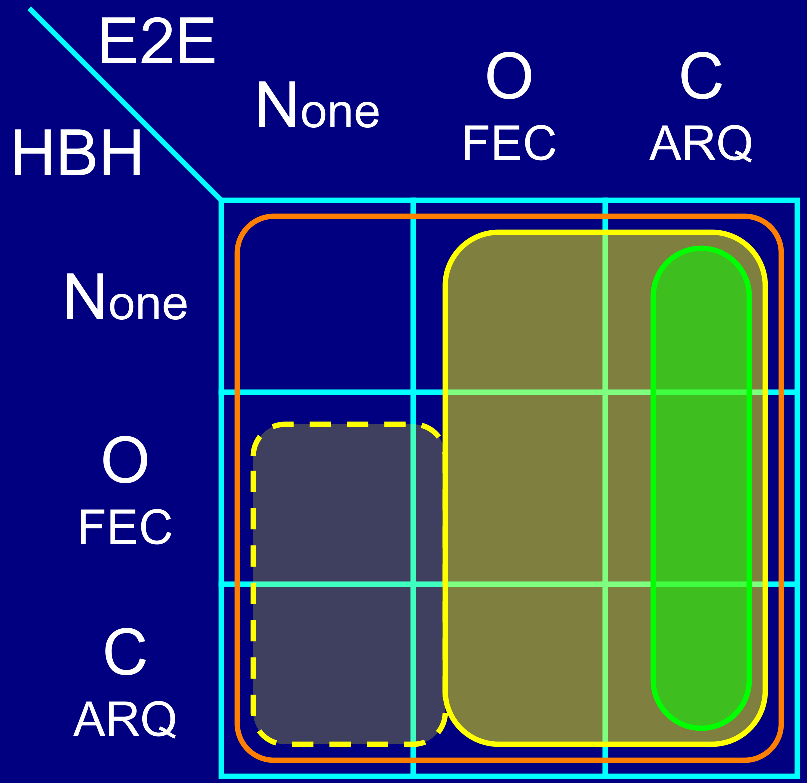

Resilient Transport: ResTP
Transmission Control

- Transmission control modes [future work]
  - subflow congestion control
  - subflows should generally not share nodes nor links

ResiliNets Protocols

Cross-Layer Model: ResTP/GeoDivRP

- Application
  - $\mathbb{K}_{7\to4} = \{\text{ss, tm}\}$
    service spec and threat model
- E2E Transport: ResTP
  - erasure spreading vs. hot standby
  - FEC vs. HARQ vs. ARQ
  - $\mathbb{K}_{4\to3} = \{k,\ d,\ [h,\ t]\}$
    k-path diversity over distance d
    opt. stretch h and skew t bounds
- Routing: GeoDivRP
  - construct k d-diverse paths

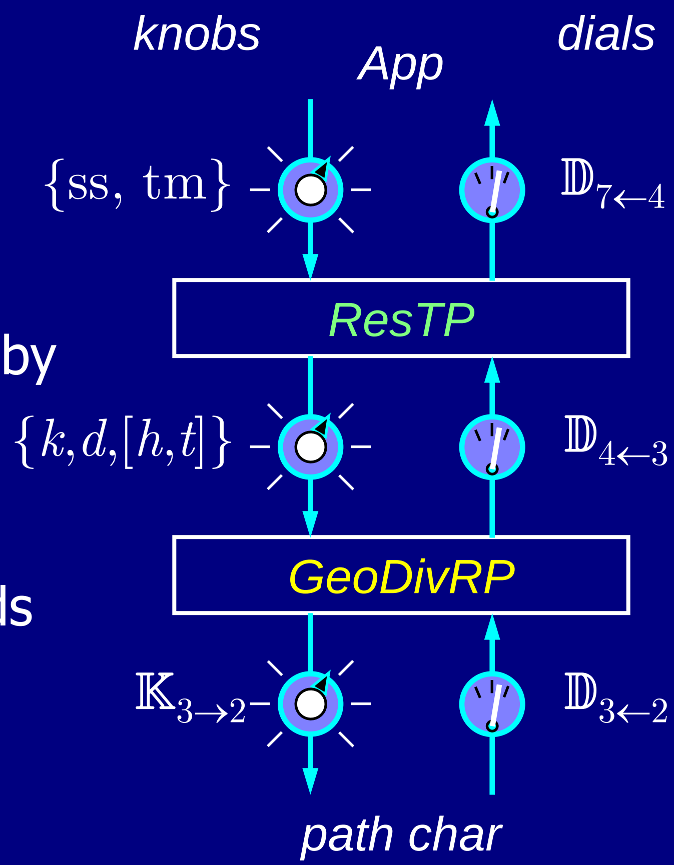

Geodiverse Routing Protocol
GeoDivRP using iWPSP and MLW

- Two heuristics: iWPSP and MLW
- iWPSP (iterative waypoint shortest path)
  - choose neighbours and waypoints to meet diversity spec
  - splice Dijkstra shortest paths
  - complexity: $2c^2n^2 \log n$ (for average of c neighbours)
  - [Cheng and Sterbenz @ KU: DRCN 2014]
- MLW (modified link weights)
  - modify link weights higher close to primary path
  - forces (weighted) shortest path alternates to be diverse
  - complexity: $2n \log n$
  - [Gardner, May, and Medhi @ UMKC: DRCN 2014]

Geodiverse Multipath Routing
GeoDivRP: iWPSP

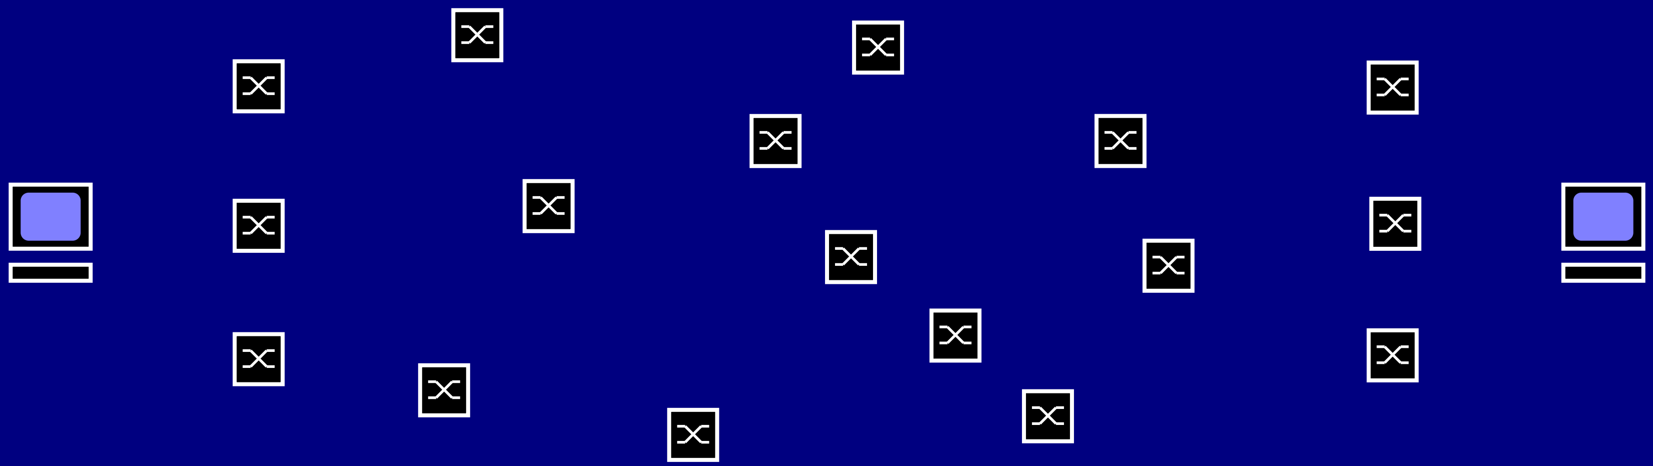

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers

Geodiverse Multipath Routing

GeoDivRP: iWPSP

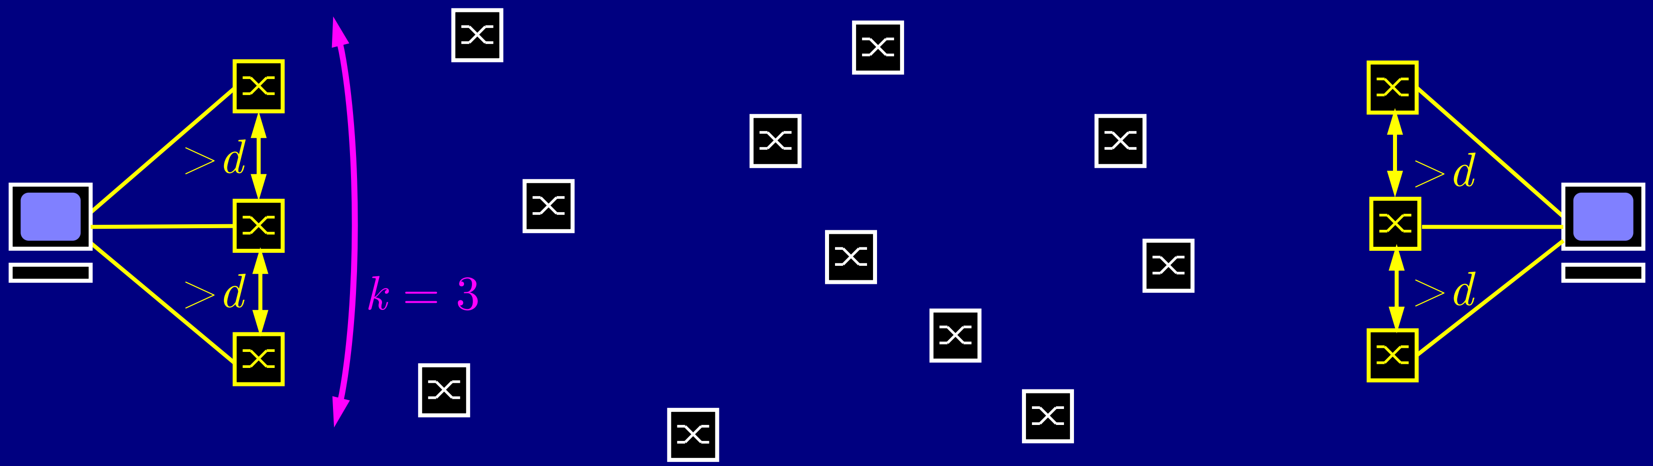

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers
  - choose k next hop routers at least d apart if possible

Geodiverse Multipath Routing
GeoDivRP: iWPSP

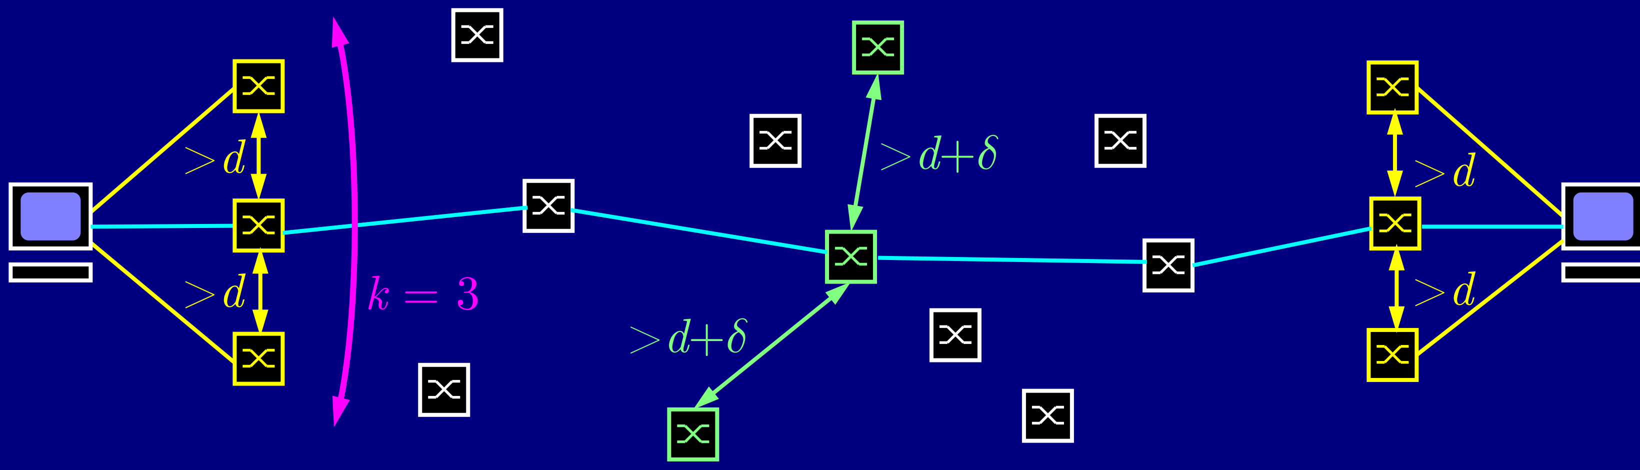

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers
  - choose k next hop routers at least d apart if possible
  - choose mid-point waypoints $d+\delta$ wrt to shortest path
    - limit stretch to h and skew to t if specified and possible

Geodiverse Multipath Routing

GeoDivRP: iWPSP

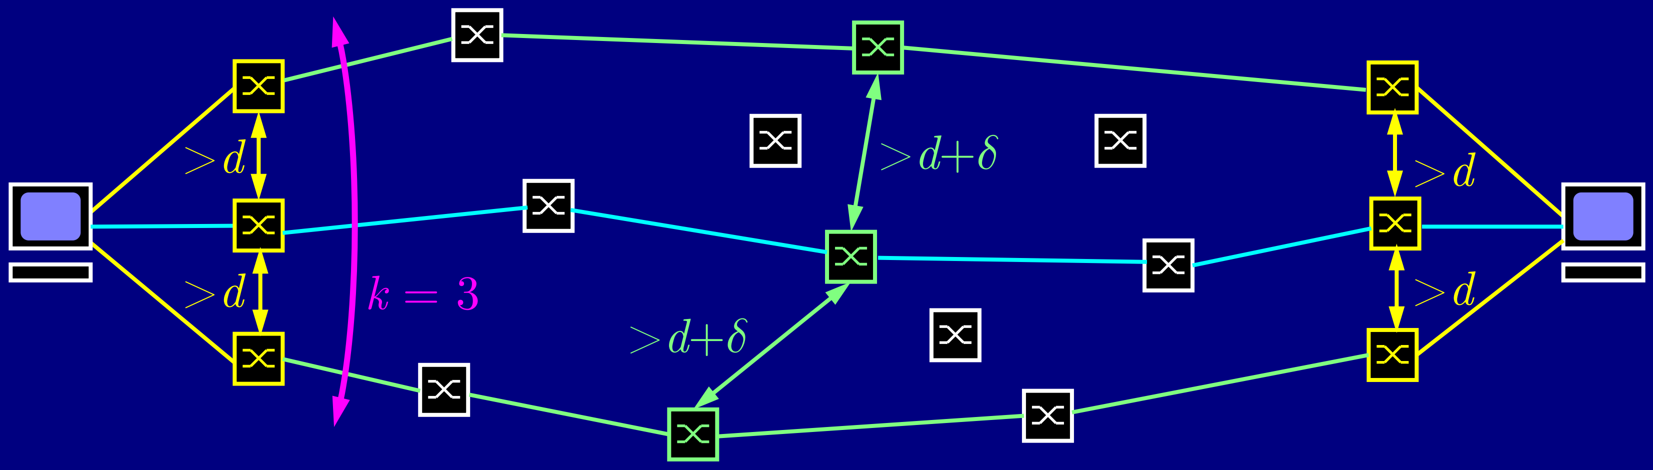

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers
  - choose k next hop routers at least d apart if possible
  - choose mid-point waypoints $d+\delta$ wrt to shortest path
    - limit stretch to h and skew to t if specified and possible
  - use conventional SPF (Dijkstra) for paths to waypoints

Geodiverse Multipath Routing

GeoDivRP: iWPSP

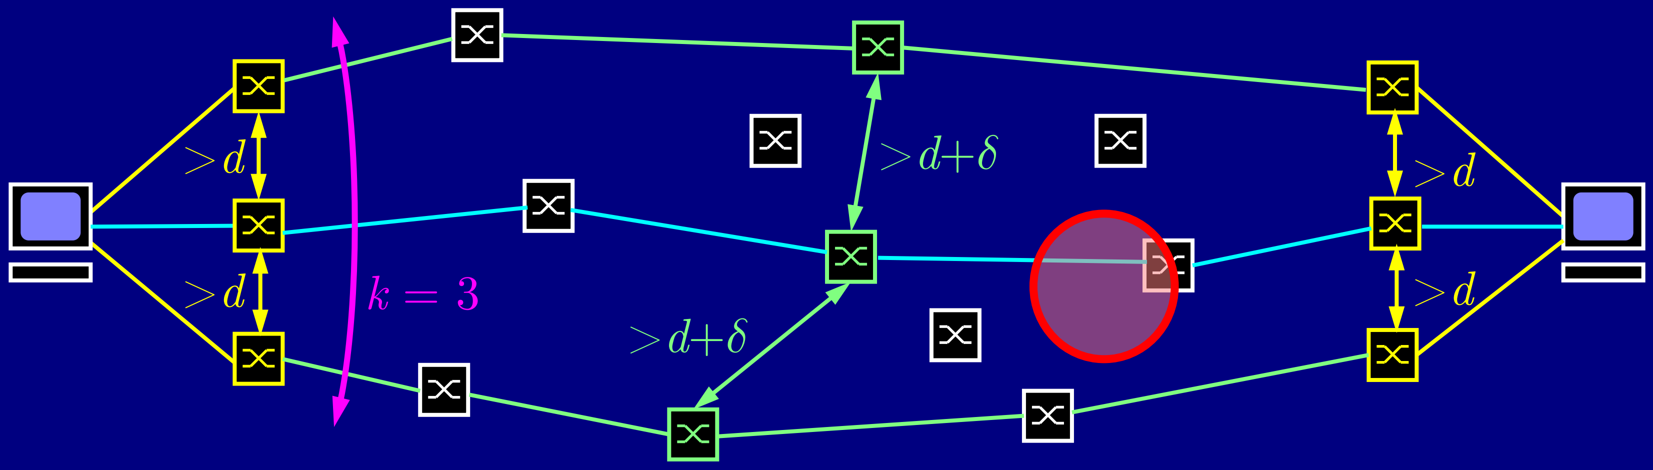

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers
  - choose k next hop routers at least d apart if possible
  - choose mid-point waypoints $d+\delta$ wrt to shortest path
    - limit stretch to h and skew to t if specified and possible
  - use conventional SPF (Dijkstra) for paths to waypoints

Geodiverse Multipath Routing

GeoDivRP: MLW

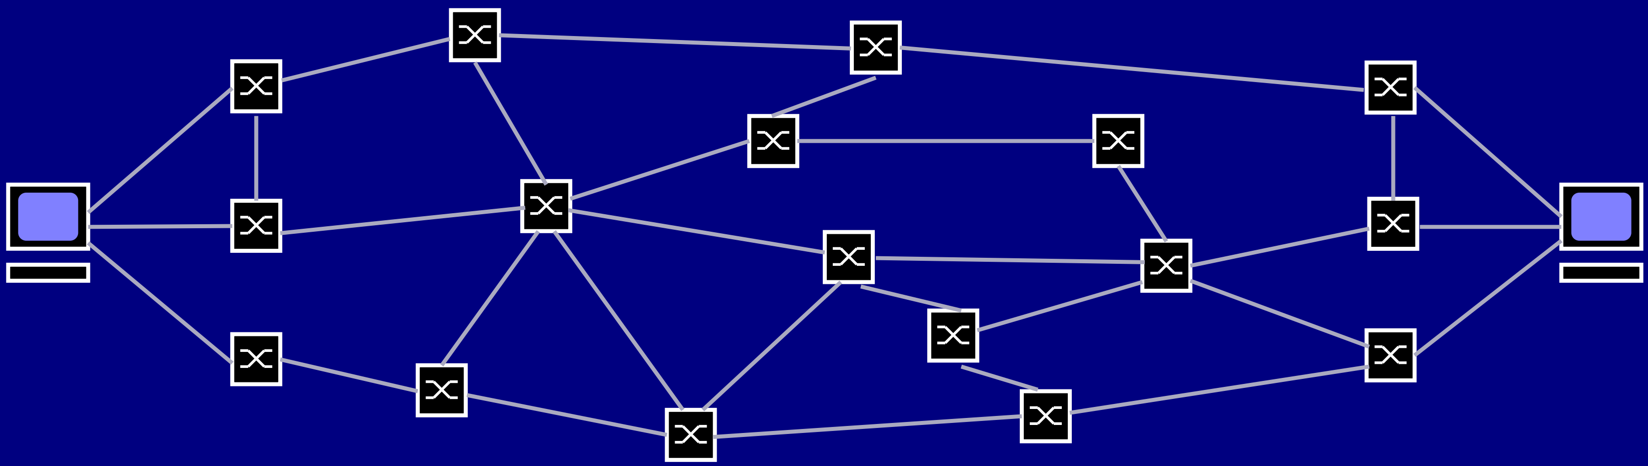

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers

Geodiverse Multipath Routing

GeoDivRP: MLW

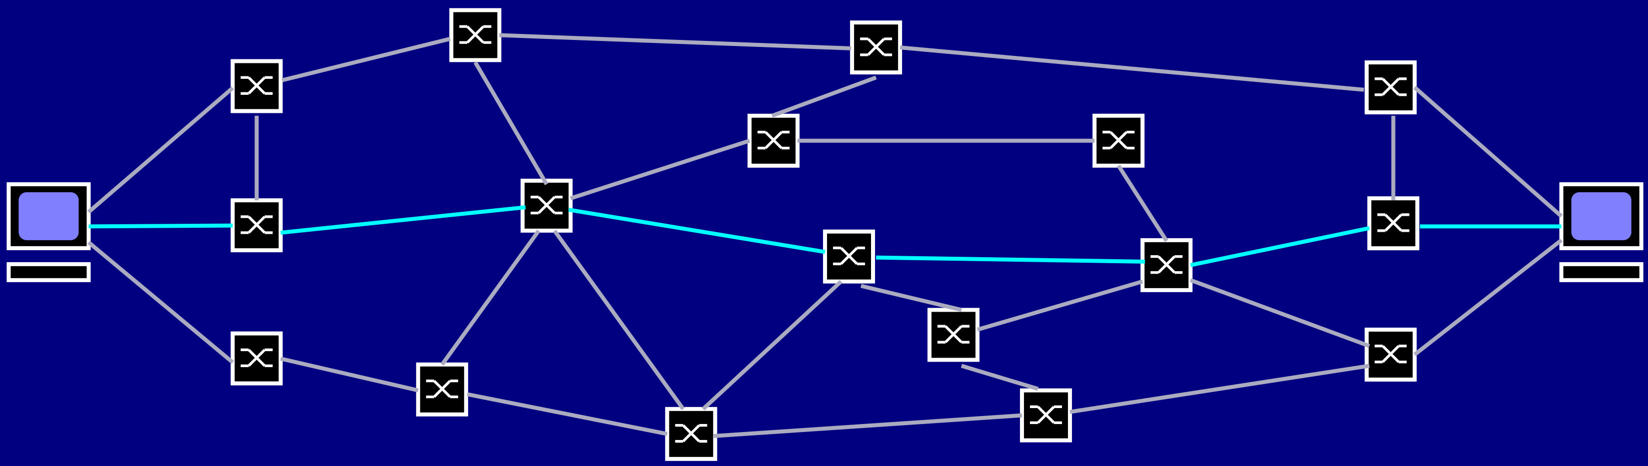

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers
  - choose primary shortest path

Geodiverse Multipath Routing

GeoDivRP: MLW

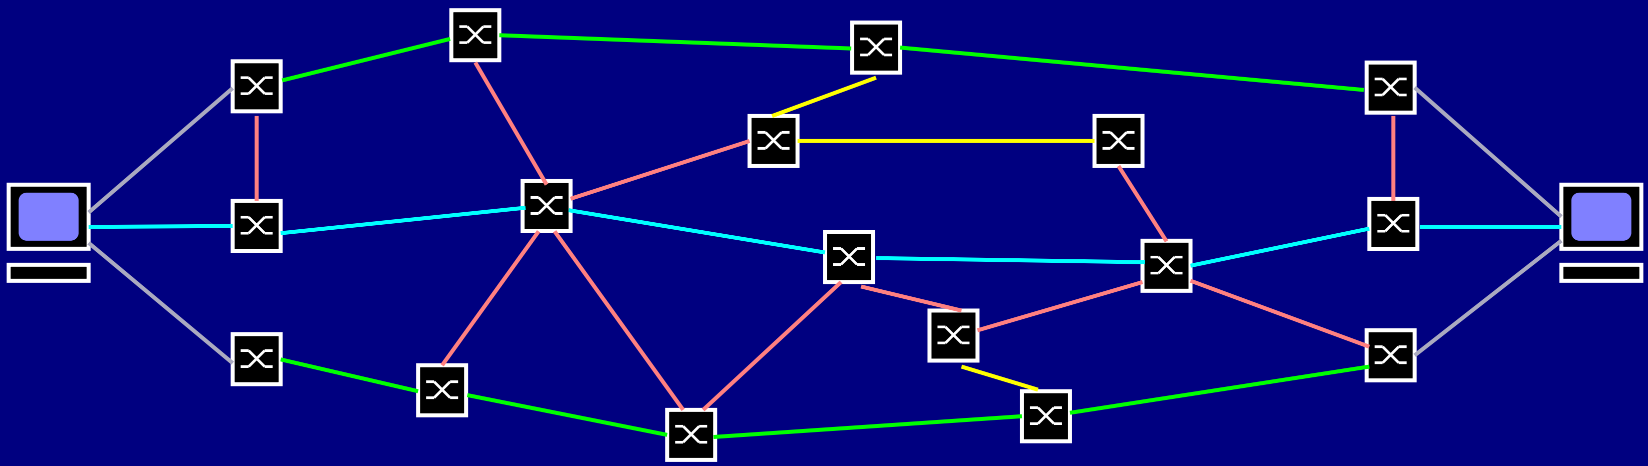

- GeoDivRP: intermediate waypoint algorithm
  - LSAs contain geolocation of routers
  - choose primary shortest path
  - modify link weights higher close to primary path
  - forces (weighted) shortest path alternates to be diverse

End