

End-to-End Communication for Resilience under Challenges in the Postmodern Internet

James P.G. Sterbenz

Abdul Jabbar, Justin P. Rohrer, Egemen Çetinkaya

Qian Shi, Rabat Mahmood

David Hutchison, Güneş Erçal-Özkaya

Department of Electrical Engineering & Computer Science
Information Technology & Telecommunications Research Center
The University of Kansas



jpgs@ittc.ku.edu

<http://www.ittc.ku.edu/~jpgs>

<http://wiki.ittc.ku.edu/resilinets>

E2E Communication in PoMo

Outline

- Motivation and introduction to PoMo
- End-to-end redundancy and diversity
- Cross-layering in PoMo
- Methodology: topology and challenge generation

Postmodern Internet

Selected Motivation


- Heterogeneity
 - Internet architecture increasingly constrains choices
 - embrace alternative technologies and mechanisms...
...giving end users choice and influence
- Tussle
 - trust and policy boundaries poorly supported
- Resilience
 - Internet relatively resilient in a global sense
 - but not locally
 - vulnerabilities can be attacked: topology, DNS, BGP, ...
 - e.g. prefix hijacking, submarine cable cuts, Baltimore tunnel fire

Postmodern Internet

Overview

- PoMo Principles

- new internetworking layer
- strict separation of concerns
- heterogeneous *realms*:
mechanism, policy, trust
- include mechanisms supporting all foreseeable policies



- Novel characteristics

- separate path determination from forwarding: source routes
- unforgeable record of path traversed by each packet
- separate customer-provider relationship from topology
- support cross-layer knobs and dials

Postmodern Internet Architecture

Header₁



- Forwarding directive (FD): *where*
 - tells infrastructure how to relay the packet
 - some combination of:
 - destination specification + partial path
 - path = sequence of channel IDs
- Motivation: *why*
 - convinces intermediate node it should relay the packet

UK

Postmodern Internet Architecture

Header₂



- **Accountability:** *who*

UMd

- System relies on a trusted identity provider (TAP)
 - no key exchange between non-neighbor ASs
 - no trusted hardware, including routers and DHT nodes
- Each packet embeds a (blinded) identity
 - packets can be traced (with TAP assist) to sending principal
- Network admins may choose to block any sending principal
 - block state held in a secure DHT

Postmodern Internet Architecture

Header₃



- **Knobs:** *how*
 - advice to network layer (and below) from above
- **Dials:** *what*
 - instrumentation from realms and inter-realm paths

KU

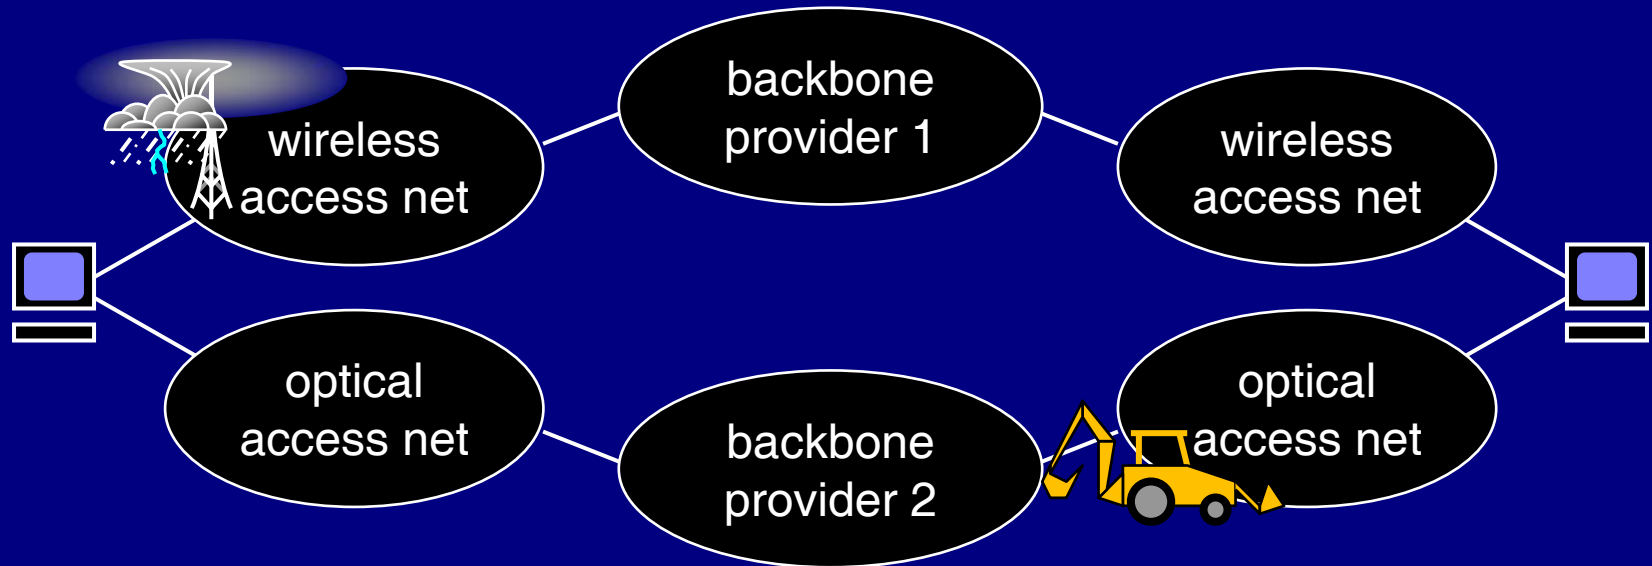
End-to-End Communication

Redundancy and Diversity

- E2E transport over multiple diverse paths
 - that have minimal (if any) shared fate
- Diversity in
 - service provider: resilience to contract and peering disputes (e.g. Cogent vs. Level3)
 - underlying technology: resilience to medium challenge
 - e.g. weather disruption of wireless mesh links
 - path geography: resilience to natural disaster and attacks
 - e.g. Baltimore tunnel fire, Hinsdale central office fire
 - fault tolerance necessary but not sufficient for survivability
- Diversity at all layers

End-to-End Communication

Example Scenario



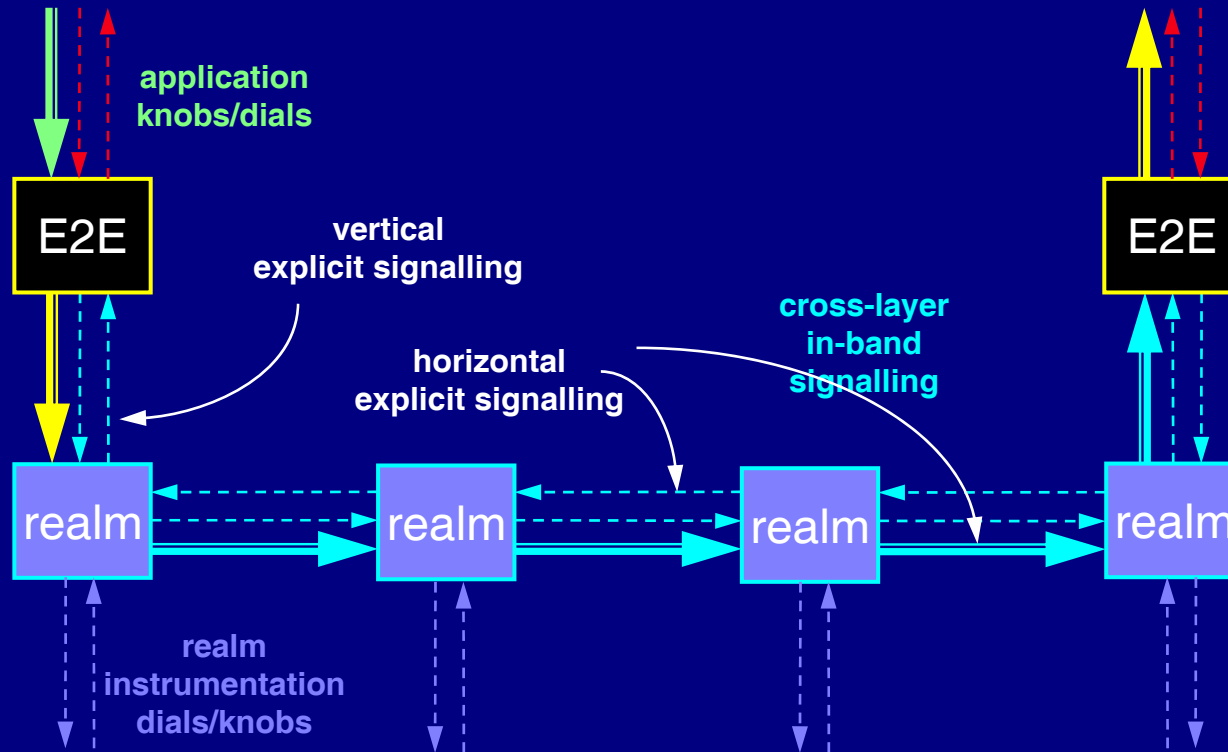
- Realm path choices explicitly available to end user
 - spreading (e.g. erasure coding) or hot standby
 - service tradeoffs: optical when available, fail-over to wireless
 - cheapest path under dynamic pricing

End-to-End Communication Requirements

- Dependent on
 - application service requirements, e.g. reliability (knobs)
 - inter-realm path diversity (dials from PoMo routing)
 - possible path source routes from PoMo routing
 - realm characteristics (dials from realm layer)
- Adaptive and composable mechanisms
 - flow type: connections, opportunistic conn., connectionless
 - error control: ARQ, ARQ+FEC, FEC, none
 - multipath options: spreading, hot-standby, failover
[TP++], [DTN], ...

End-to-End Communication

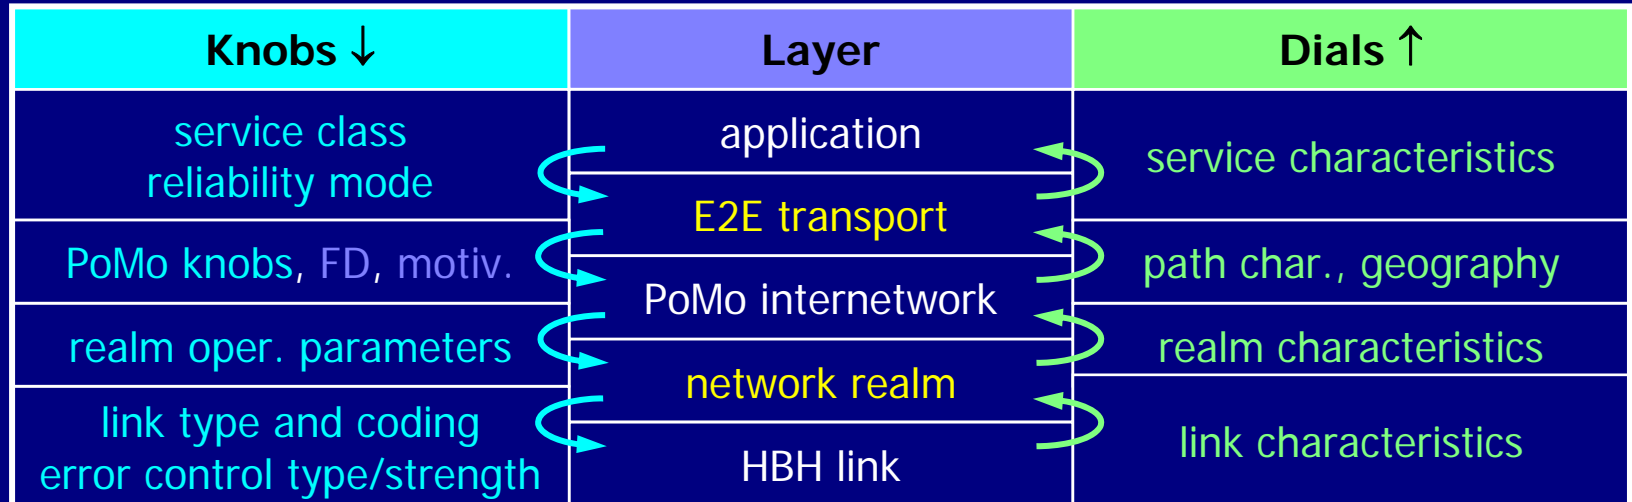
Knobs and Dials in PoMo Header



- In band: knobs and dials in data packet
- Explicit signalling: knobs and dials in signalling packet

End-to-End Communication

Knobs and Dials



- Knobs and dials between upper layers and PoMo
 - support heterogeneous subnetworks
 - e.g. lossy wireless vs. reliable wired
 - explicit signalling of path diversity and multipath
 - geographic location of realms, nodes, channels

End-to-End Communication

Reliability Mode Example

- Reliability modes:
 - reliable: ARQ with E2E ACKs
 - near-reliable: custody transfer from realm-to-realm
 - quasi-reliable: open-loop FEC or erasure coding
 - unreliable: relies on lower layer best effort service
- Multipath with spatial diversity for reliability

End-to-End Communication


Role of Overlays?

- Multilevel abstractions
 - cross-layer interactions a difficult problem
 - realm characteristics propagates up to application
- Overlays make this much harder: obscure diversity
 - are overlays a desired outcome?
or
 - a necessary evil of the current Internet architecture

Evaluation Methodology

Flexible and Realistic Topology Generation



- Hierarchical topology generation
 - evaluation of PoMo mechanisms
 - network engineering for resilience
- Level 1: **backbone realms**
 - nodes distributed based on location constraints
 - links generated using various models under cost constraints
- Level 2: **access network realms**
 - distributed around backbone nodes
 - access network connectivity: ring, star, mesh
- Level 3: **subscribers**
 - distributed around access network node



Evaluation Methodology

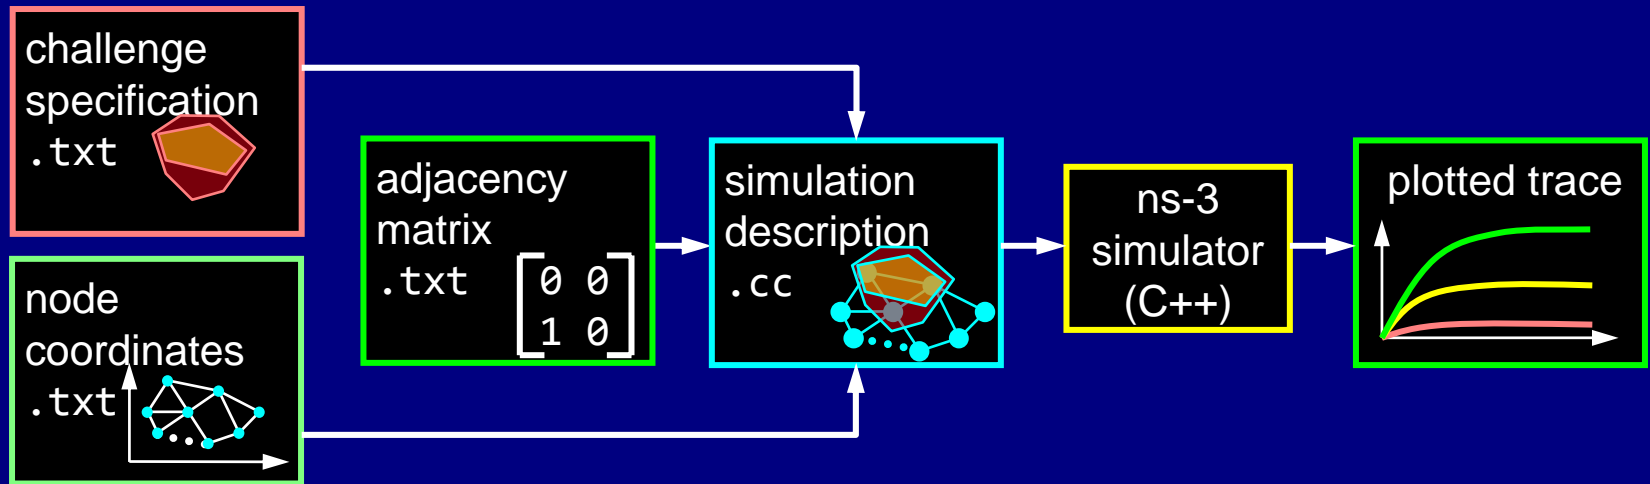
Challenge Simulation Module

- Separate challenge from network simulation
- Simulate challenges to any network over time interval
 - natural disaster: polygon destroys network infrastructure
 - attack: {node|link} down, wireless link attenuated



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

KU-LoCGen

Evaluation Methodology

Resilience Trajectory

- Resilience

- trajectory through 2-dimensional state space
- operational dimension given challenges

- normal operation
- partially degraded
- severely degraded

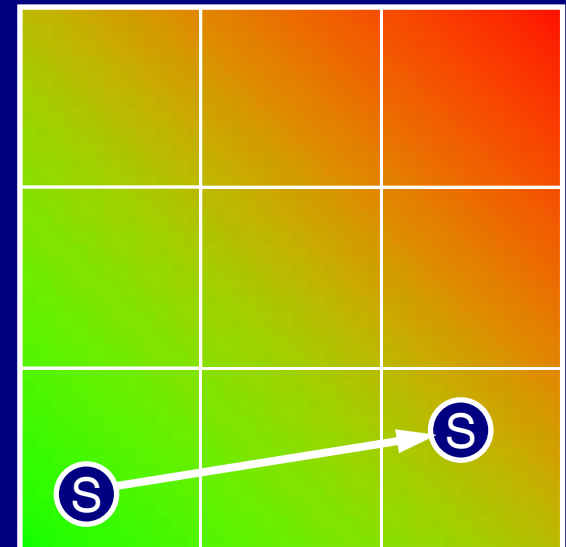
- service dimension

- acceptable service
- impaired but usable service
- unacceptable service

Unacceptable

Impaired

Acceptable



normal operation partially degraded severely degraded

PoMo E2E Communication

Implementation Status

- KU-LoCGen and KU-CSM near completion
 - MATLAB and ns-3
- PoMo end-to-end transport
 - preliminary analysis done in ns-2
 - simulation and cross-layering being converted to ns-3
 - slow and painful process that will hopefully pay of in long run
 - cross-layering and heterogeneity required ns-2 source hacks
 - implementation planned over UK forwarding code
 - deployment and testing planned in GpENI

End