# KIRA – Scalable Zero-Touch Routing for Autonomic Control Planes

Roland Bless, Martina Zitterbart
Institute of Telematics, KIT

Zoran Despotovic, Artur Hecker
Huawei Research Center, Munich

R. Bless, M. Zitterbart, Z. Despotovic and A. Hecker, "KIRA: Distributed Scalable ID-based Routing with Fast Forwarding„, 2022 IFIP Networking Conference (IFIP Networking), 2022, https://s.kit.edu/KIRA

# Internet Infrastructure…

- is becoming more complex
  - higher interdependencies of services
- must be reliable → resilient operation

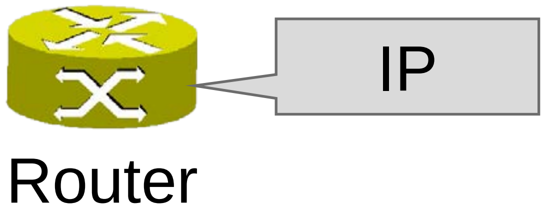

# Internet Infrastructure…

- is becoming more complex
  - higher interdependencies of services
- must be reliable → resilient operation

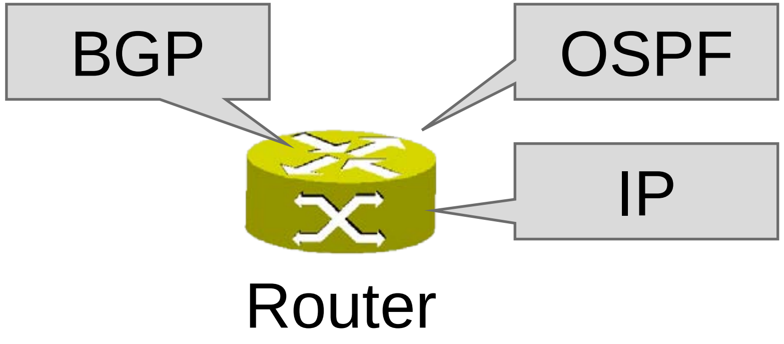

# Internet Infrastructure…

- is becoming more complex
  - higher interdependencies of services
- must be reliable → resilient operation

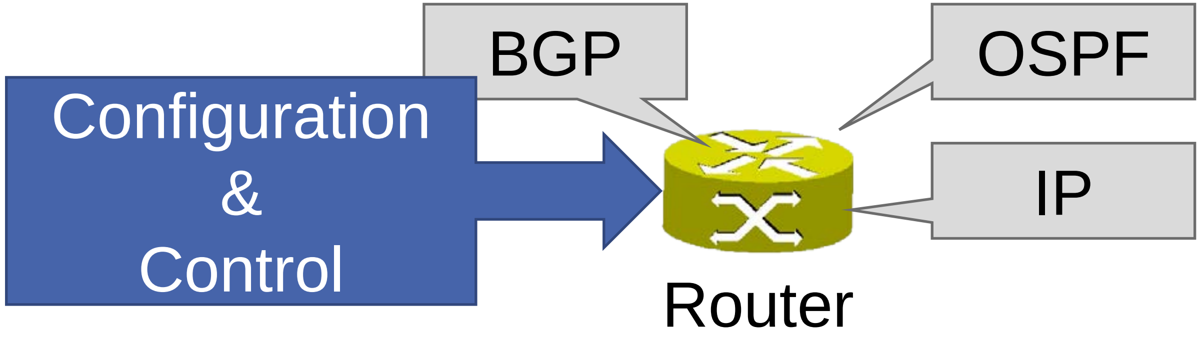


- Requires configuration via management/control plane

**This Talk**

**Foundation for Resilient Internet Infrastructures**

# Controllability and Control Planes

**Major Service Disruptions**

Meta

Oct 2021

aws

Dec 2021

# Controllability and Control Planes

Karlsruhe Institute of Technology

**Major Service Disruptions**

Meta

Oct 2021

aws

Dec 2021

Services depend on resilient connectivity

**Control plane** connectivity inherently important

# Controllability and Control Planes

**Major Service Disruptions**

- Meta – Oct 2021
- aws – Dec 2021

→

Services depend on resilient connectivity

**Control plane** connectivity inherently important

**KIRA** → provides **self-organized robust** control plane connectivity

# Control Planes of Future Networks Need to Support…

Karlsruhe Institute of Technology

Interconnection of a Large Pool
of Networked Resources

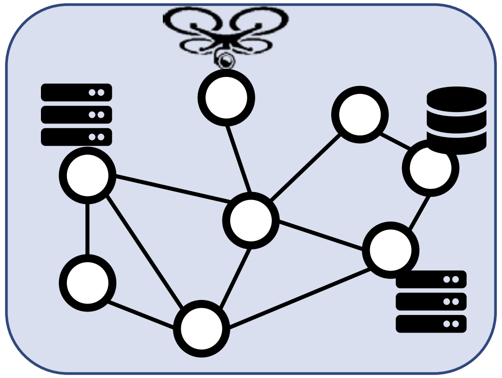

Compute, Storage, Network

- Scalability
- In-band control
- High dynamics
- Multiple domains
- Various topologies

Resilient Connectivity
for Control Plane

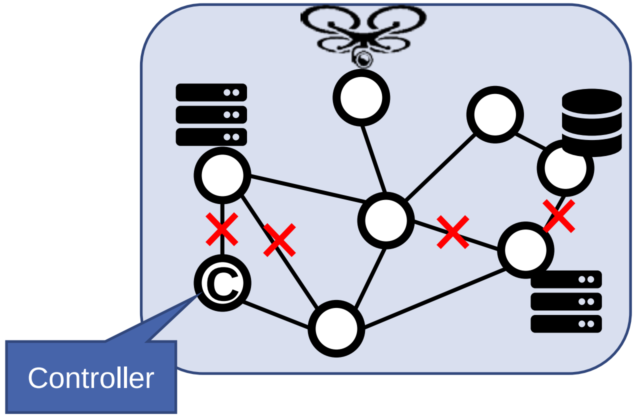


- Zero-touch
- Fast convergence
- Network split
- Nomadic networks

Stable Addresses
for Moving Resources

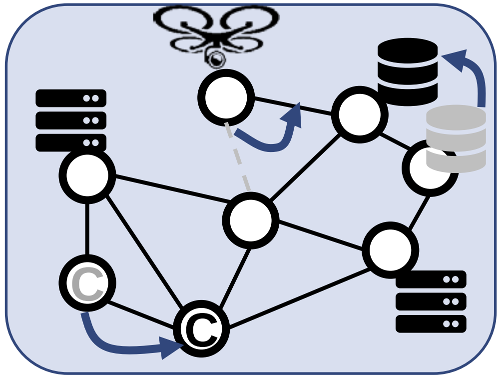


- ID-based addresses

# What KIRA achieves…

Interconnection of a Large Pool of Networked Resources

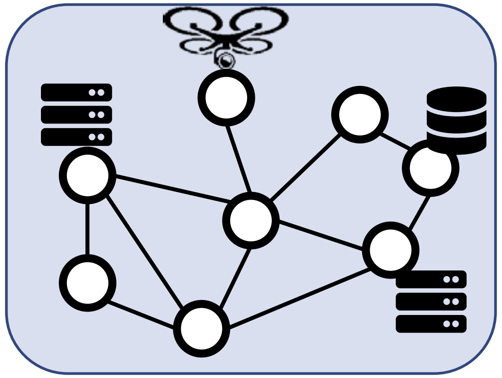

Compute, Storage, Network

Resilient Connectivity for Control Plane

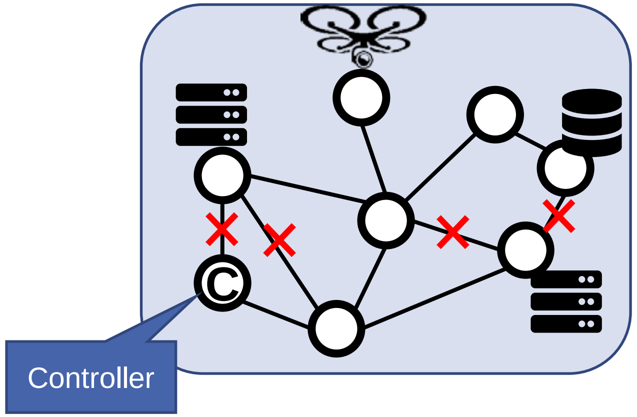


Stable Addresses for Moving Resources

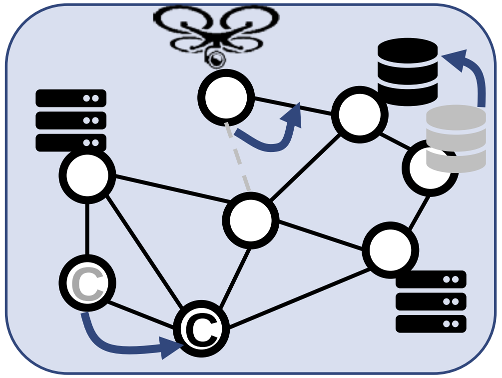


- KIRA provides (all-in-one)
  - Massive scalability (100,000s of nodes)
  - Zero-touch (no configuration)
  - Dynamics: fast convergence, loop free
  - Topological versatility
  - Efficient routes
- Related Works (examples)
  - UIP: lacks dynamics and efficient routes
  - DISCO: lacks dynamics
  - RIFT, Data Center BGP/OSPF/IS-IS: specific topologies only, not ID-based
  - RPL: traffic concentration near root, zero-touch?
  - …

# What KIRA provides…

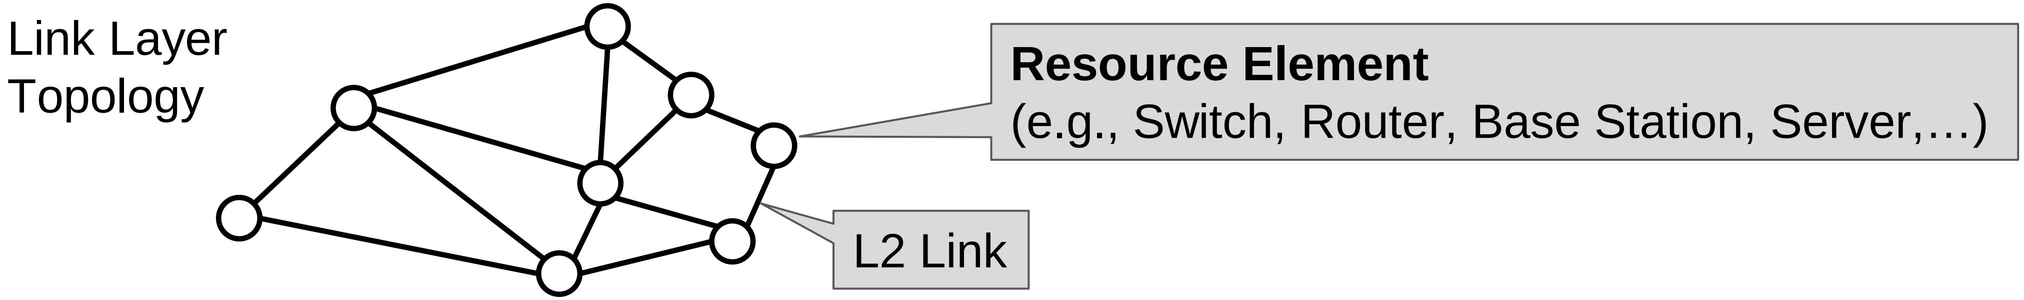

# What KIRA provides…

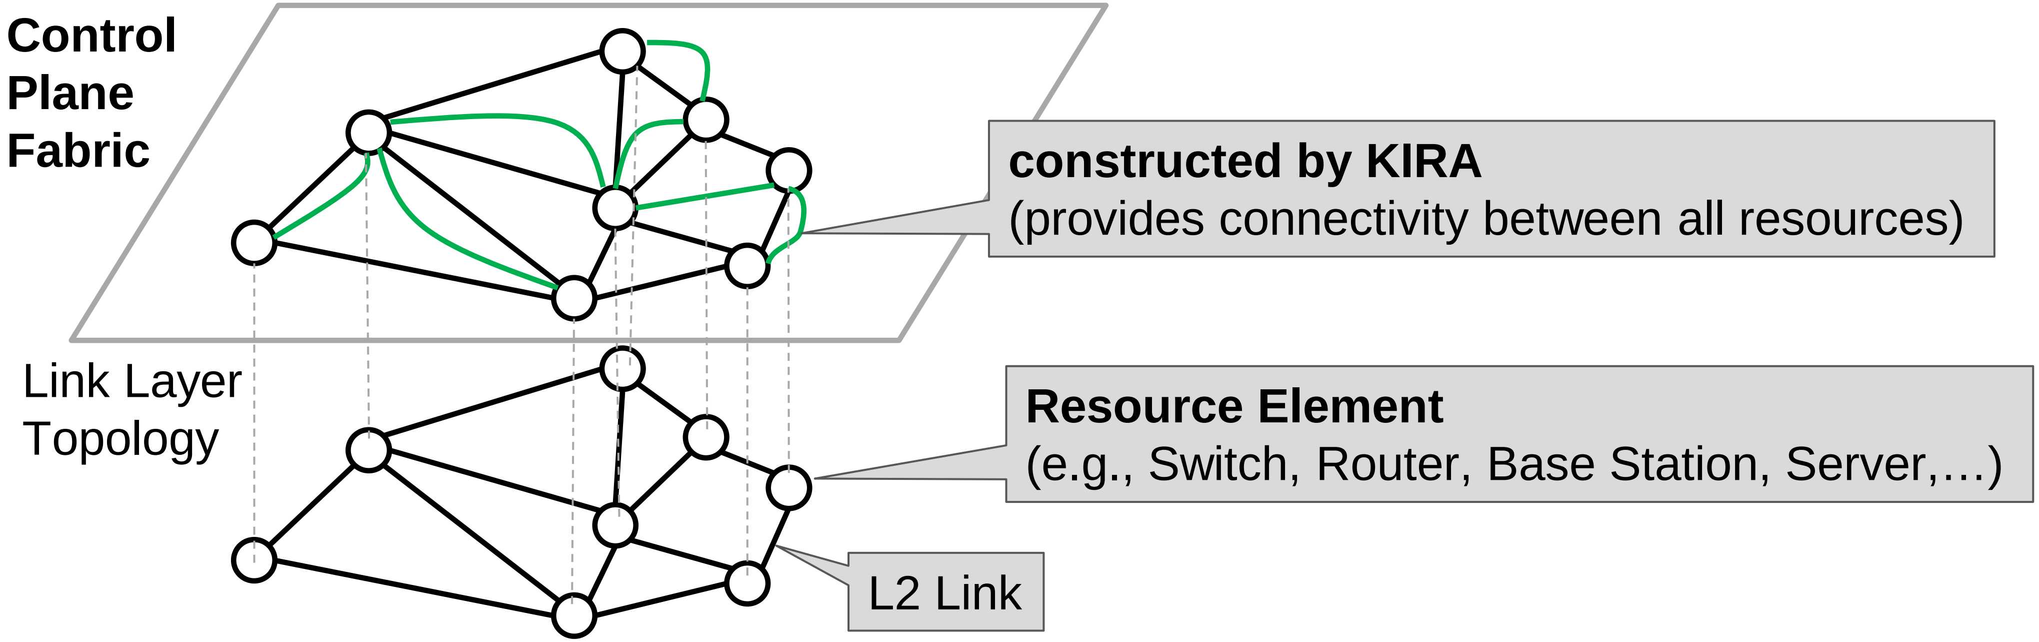

# What KIRA provides…

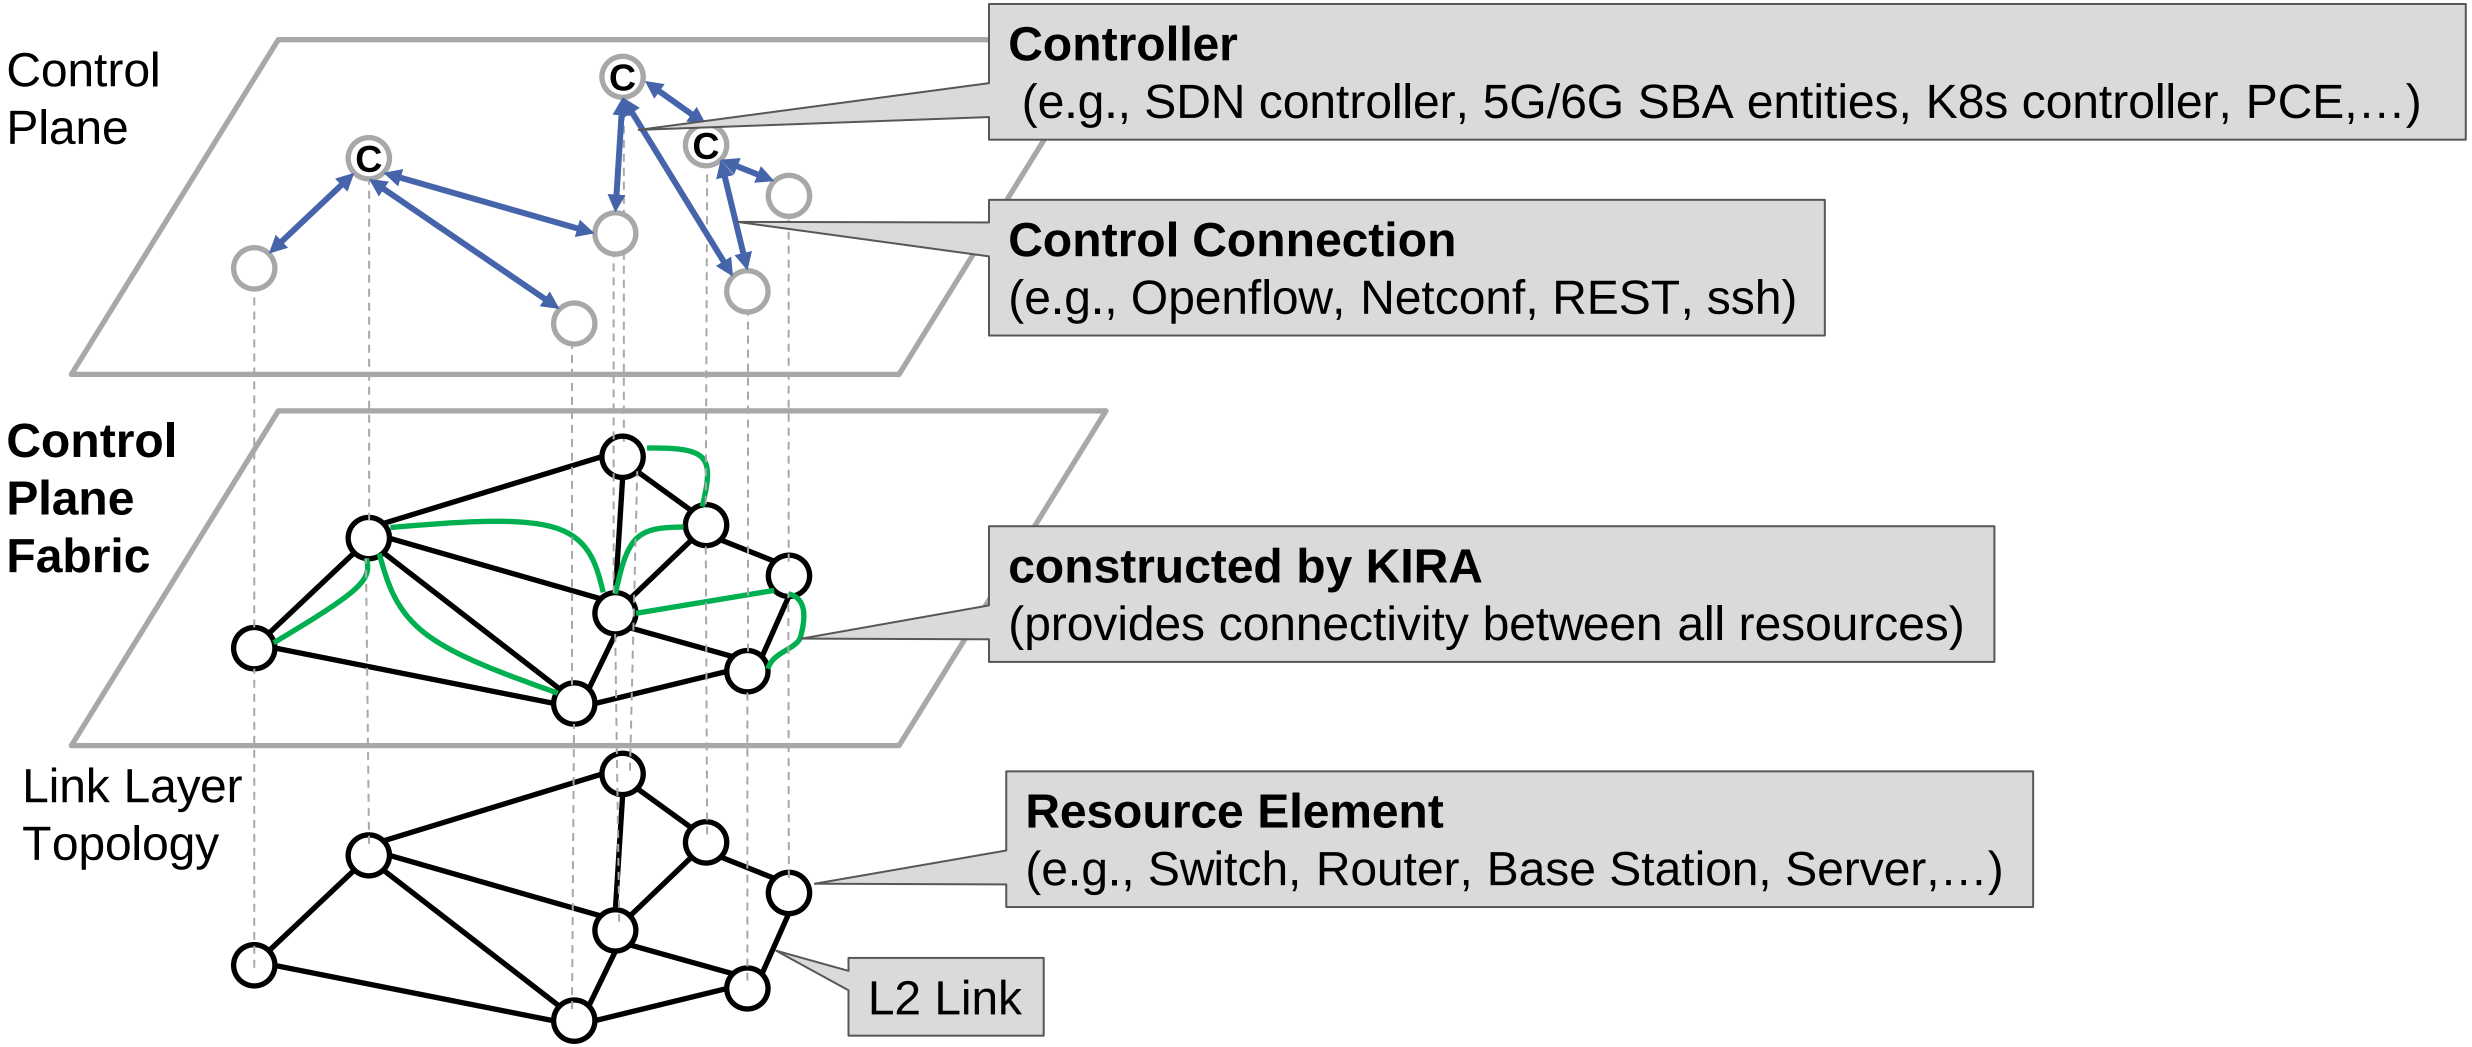

# What KIRA provides…

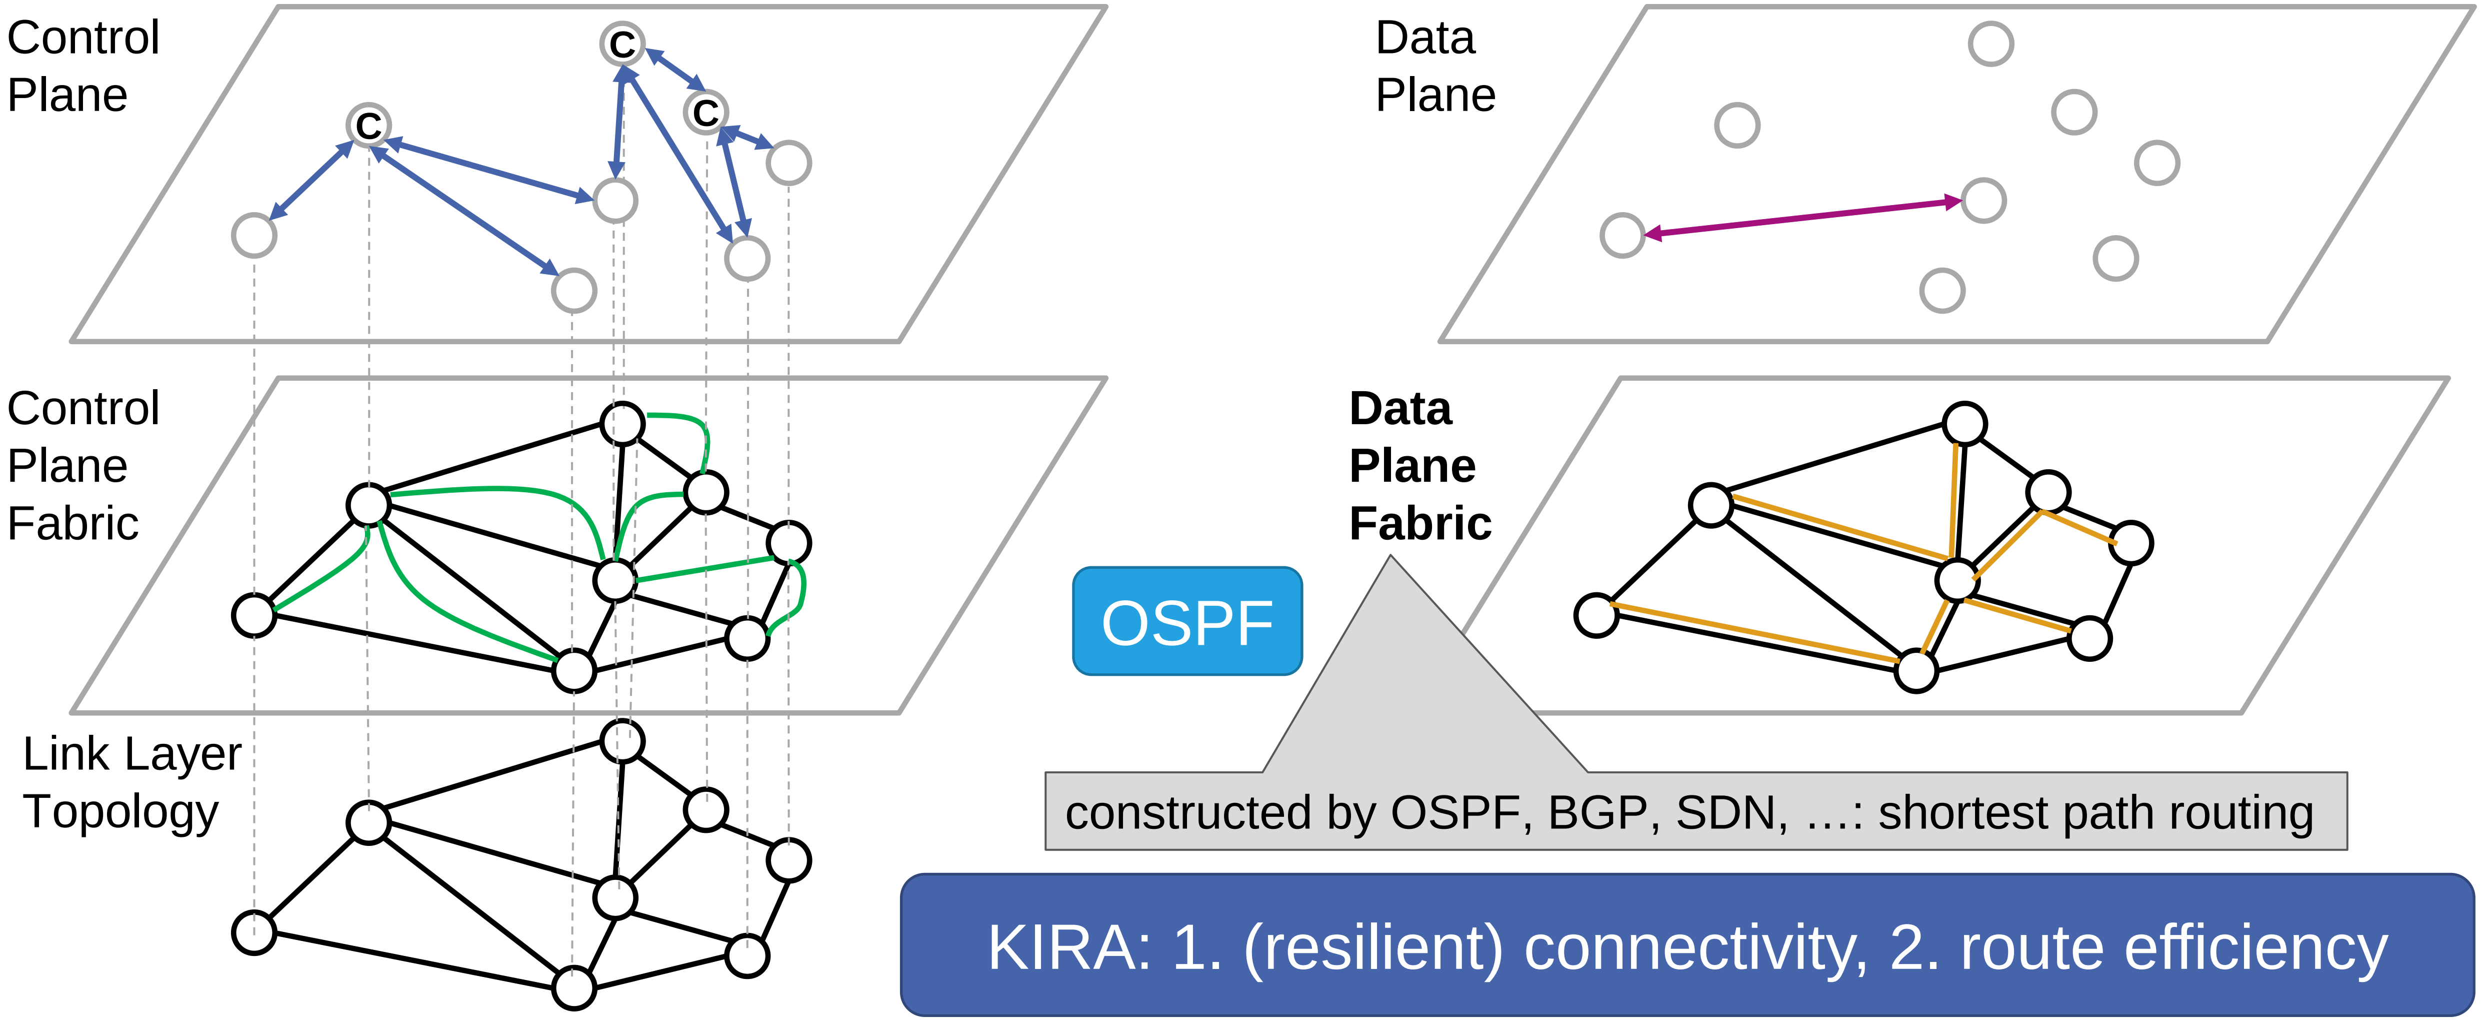


KIRA: 1. (resilient) connectivity, 2. route efficiency

# KIRA – Main Components

- Routing Tier → connectivity

**R²/Kad**

- **ID-based addresses**
- **Source routing**
- **On top of link layer**

- Forwarding Tier → optimization

**PathID-based Forwarding**

- **Eliminates source routing**
- **Label switching approach**
- **Reduces overhead**

# KIRA – Main Components

- Routing Tier → connectivity

Path Discovery, Routing

**R²/Kad**

Failure Recovery

- **ID-based addresses**
- **Source routing**
- **On top of link layer**

- Forwarding Tier → optimization

**PathID-based Forwarding**

- **Eliminates source routing**
- **Label switching approach**
- **Reduces overhead**

# R²/Kad – Path Discovery


- Each node
  - randomly chooses its NodeID (Overlay)
  - explores its 2-hop vicinity (Underlay)
    - X learns contacts A, Y, Q, B, M, …

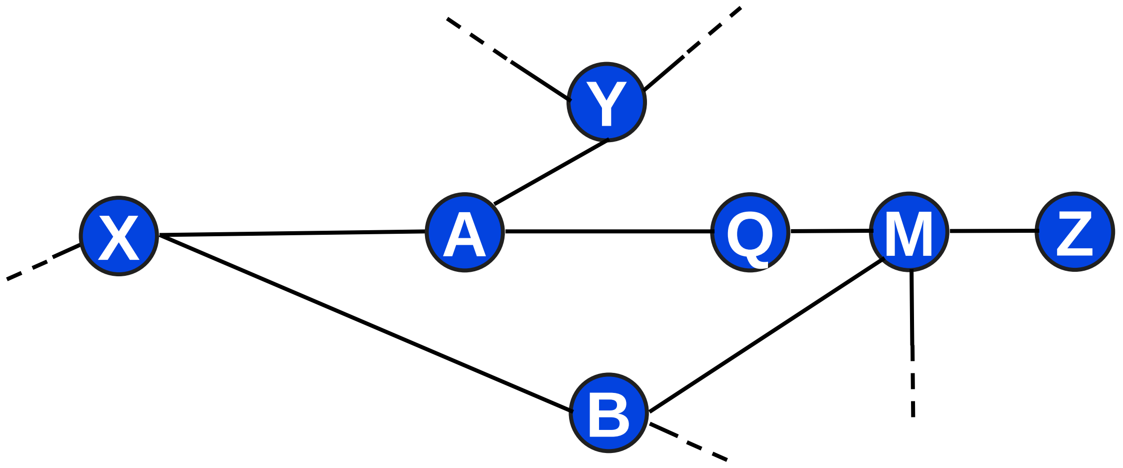

# R²/Kad – Path Discovery


- Each node
  - randomly chooses its NodeID (Overlay)
  - explores its 2-hop vicinity (Underlay)
    - X learns contacts A, Y, Q, B, M, …
- X: path to Z?
- Approach:
  construct underlay routes
  by using the NodeID-based overlay
  - Source route to contact that is
    ID-wise closest to destination NodeID (→ recursively)
  - Distance of NodeIDs: XOR metric $d(X,Y) = X \oplus Y$
    - Longer shared prefix → closer

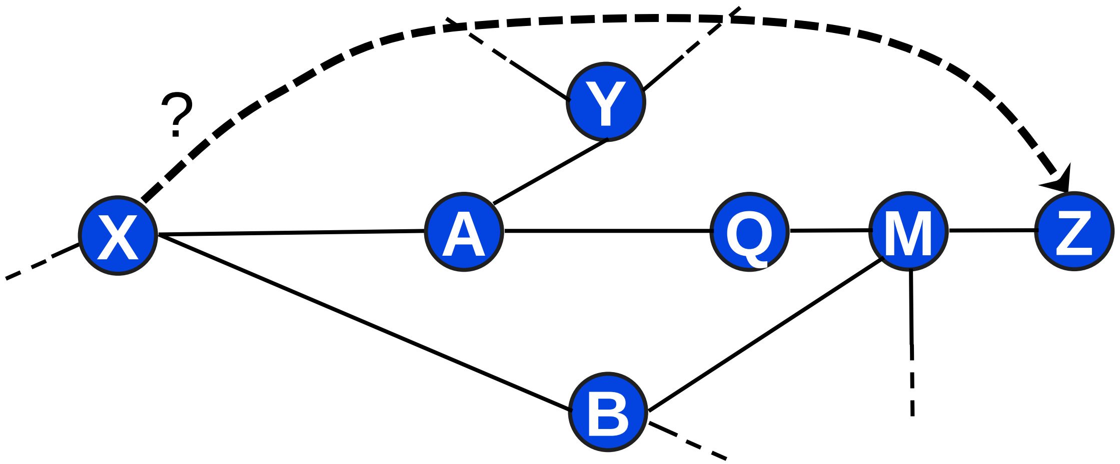

# R²/Kad – Path Discovery Example

Path Discovery, Routing | R²/Kad | Failure Recovery

- X sends FindNodeReq to contact closest to NodeID Z
  - Example: letters close in alphabet ↔ NodeIDs close
  - Next (overlay) hop: Y
- X → Y via source route <A>
- Assume Y knows Z already
- Y → Z via source route <A,Q,M>
- FindNodeReq records complete route <X,A,Y,A,Q,M>
  - incurs path stretch: $\frac{|\text{selected path}|}{|\text{shortest path}|}$

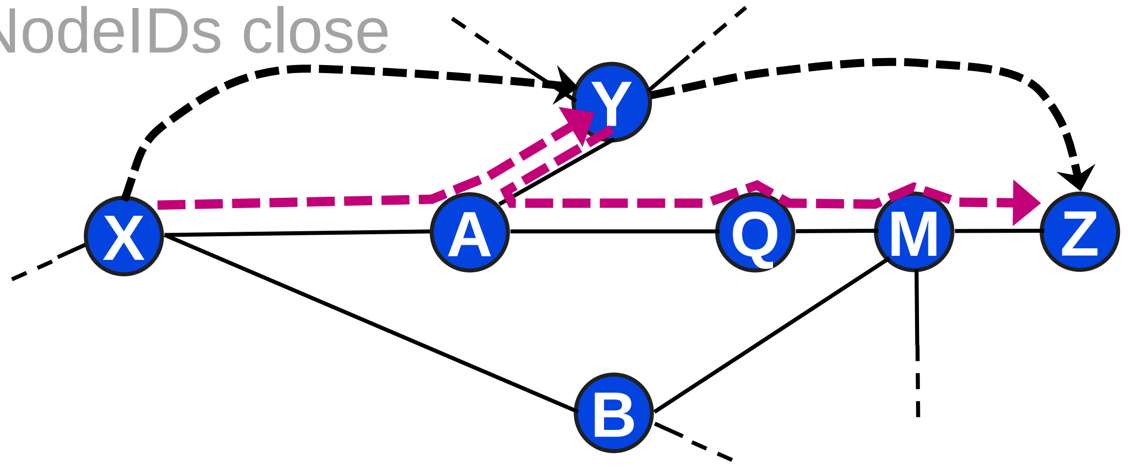

# R²/Kad – Path Discovery Example

Path Discovery, Routing | R²/Kad | Failure Recovery

- Shortened recorded route <A,Q,M> is returned to X in FindNodeRsp
- Later packets use shorter route <B,M>
  - if X already knows M via <B>

Initial stretch can be reduced for later packets!

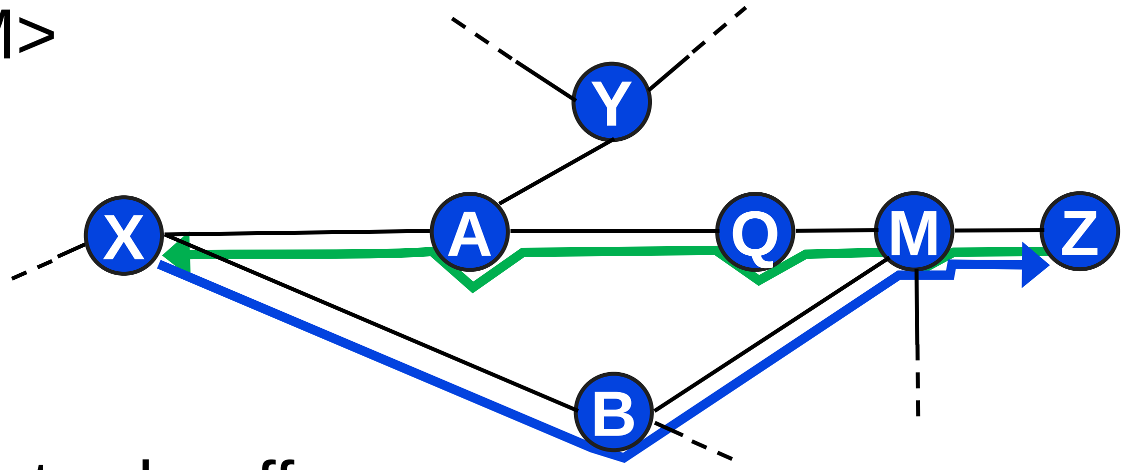


- R²/Kad offers flexible memory/stretch trade-off…

# R²/Kad – Dynamics: Rediscovery Procedure

- Detection of node/link failure in the underlay
- Two step strategy
  1.) inform ID-wise neighbors about failed link
  2.) …

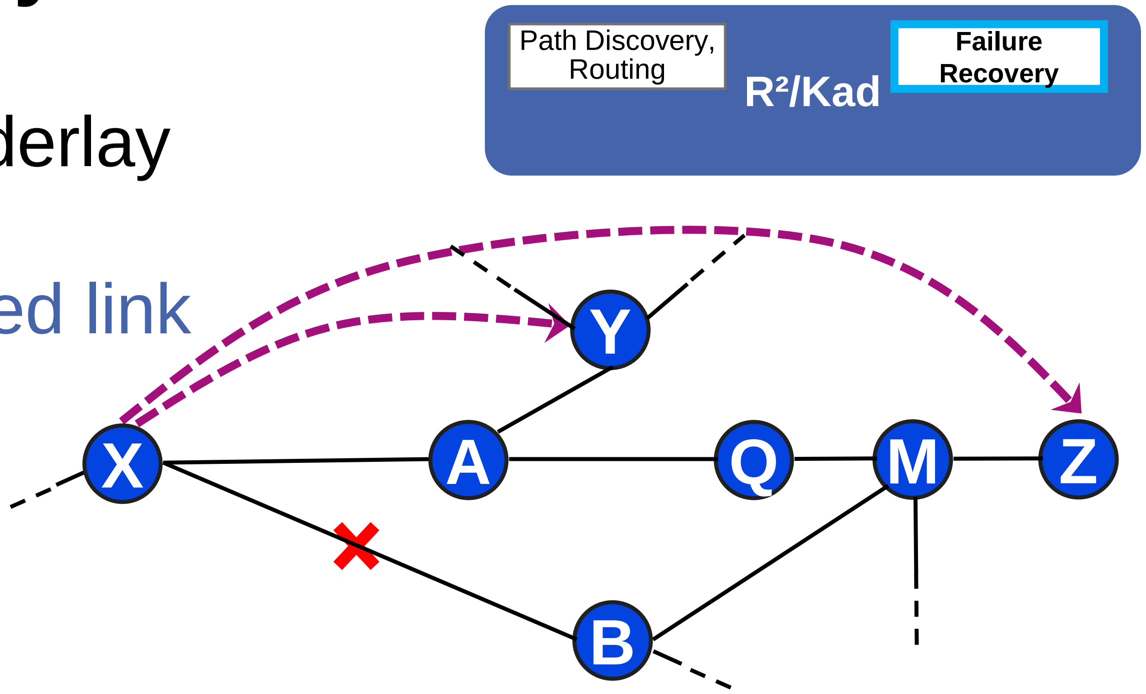

# R²/Kad – Dynamics: Rediscovery Procedure


- Detection of node/link failure in the underlay
- Two step strategy
  1.) inform ID-wise neighbors about failed link
  2.) rediscover alternative paths via overlay routes (includes "Not Via" information)

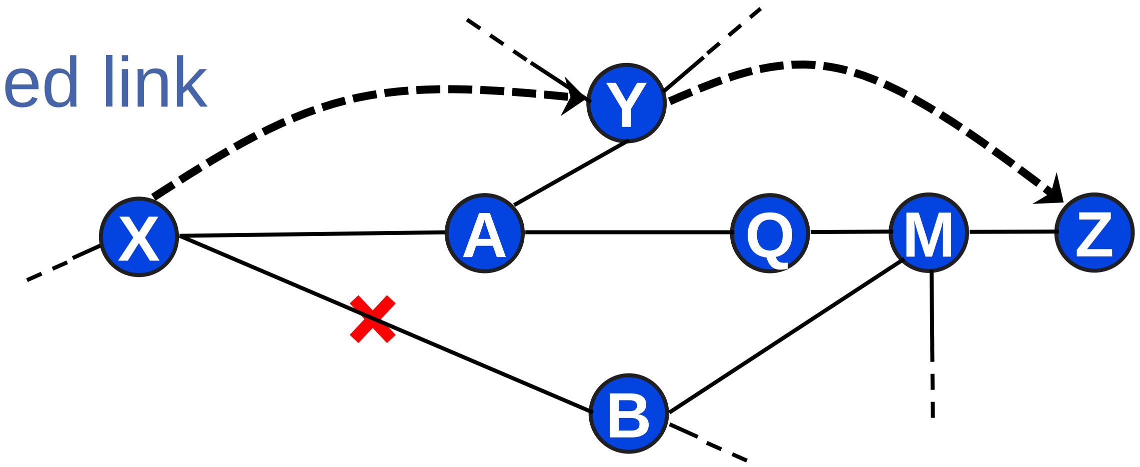


- Validity
  - State sequence numbers
  - Path information age
- Periodically
  - probe contacts for broken paths
  - lookup own NodeID

# KIRA – Main Components

- Routing Tier → connectivity

Path Discovery, Routing

Routing Table

**R²/Kad**

Failure Recovery

- **ID-based addresses**
- **Source routing**
- **On top of link layer**

- Forwarding Tier → optimization

**PathID-based Forwarding**

- **Label Switching Approach**
- **Eliminates Source Routing**
- **Reduces Overhead**

# Forwarding Tier – Fast Forwarding

- Get rid of source routes for control plane traffic
  - Reduce per packet overhead
- Approach: replace source routes with PathIDs
  - PathID(<A,Q,M,Z>)= Hash(A | Q | M | Z)
- Use PathID as unique label for path segment → Label Switching

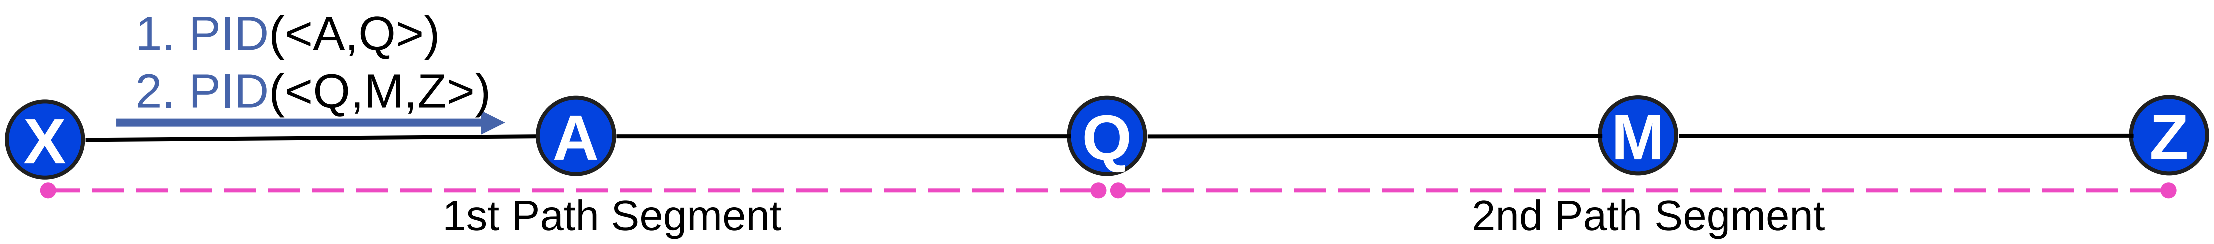

# Forwarding Tier – Fast Forwarding

- Get rid of source routes for control plane traffic
  - Reduce per packet overhead
- Approach: replace source routes with PathIDs
  - PathID(<A,Q,M,Z>)= Hash(A | Q | M | Z)
- Use PathID as unique label for path segment → Label Switching

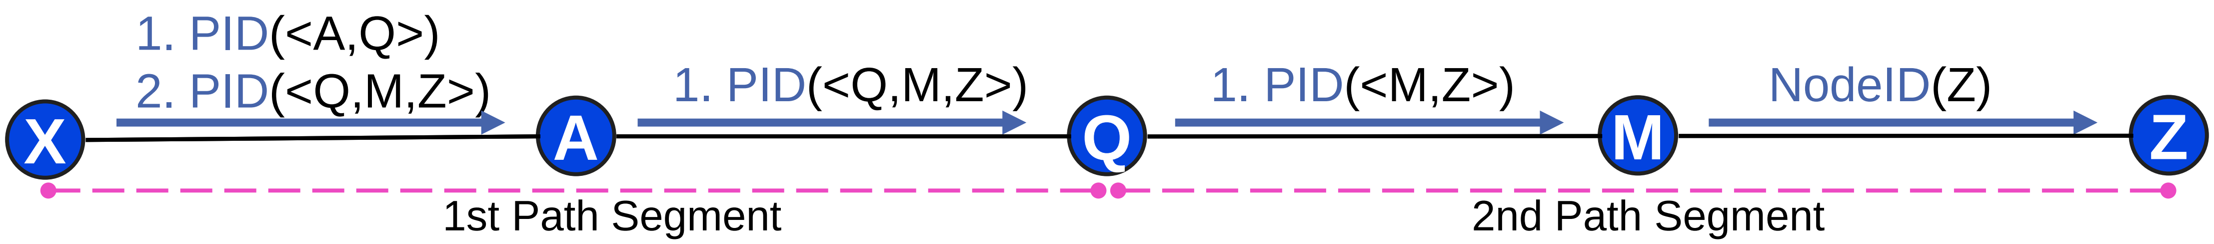


- Precalculate PathIDs for 2-hop (physical) vicinity
- Explicit path setup for paths ≥ 6 hops

# Evaluation – Simulation Setup

Karlsruhe Institute of Technology

- Simulations using RoutingSim → Dynamics (node/link failures)
- OMNeT++ 5.7
- 10 repetitions with different seeds
- Random processing time per node uniformly drawn from [0…500]µs

- Various topologies of different sizes up to 200,000 nodes:
  - Small World: Power-Law, Watts Strogatz, Internet-AS level
  - Regular: Grid, Fat Tree, Mixed Fat Tree/Power Law
  - Random: Random, Random Geometric
  - Real: Topology Zoo

# Evaluation – Topological Versatility

- Multiplicative Stretch
- Bucket size k=40
- RPL-ACP:
  - Storing-mode
  - Single DODAG
  - Single DODAG version

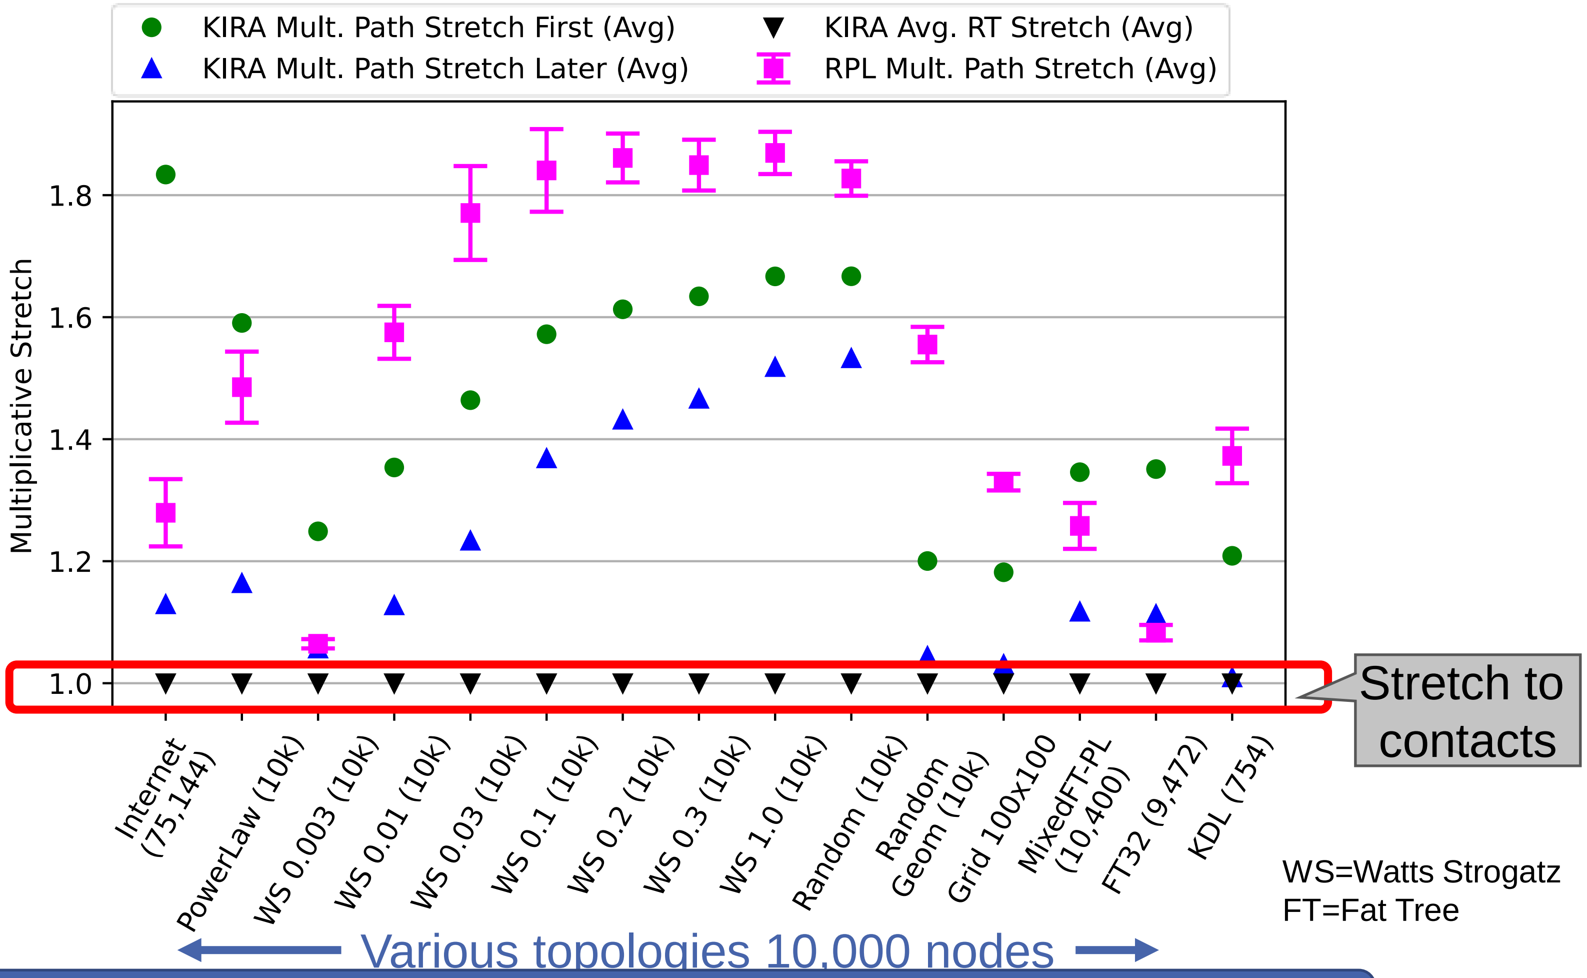


WS=Watts Strogatz
FT=Fat Tree

**Low stretch across various topologies + Shortest paths to contacts**

# Evaluation – Dynamics

- 100.000 nodes Power-Law

15% links fail randomly and simultaneously

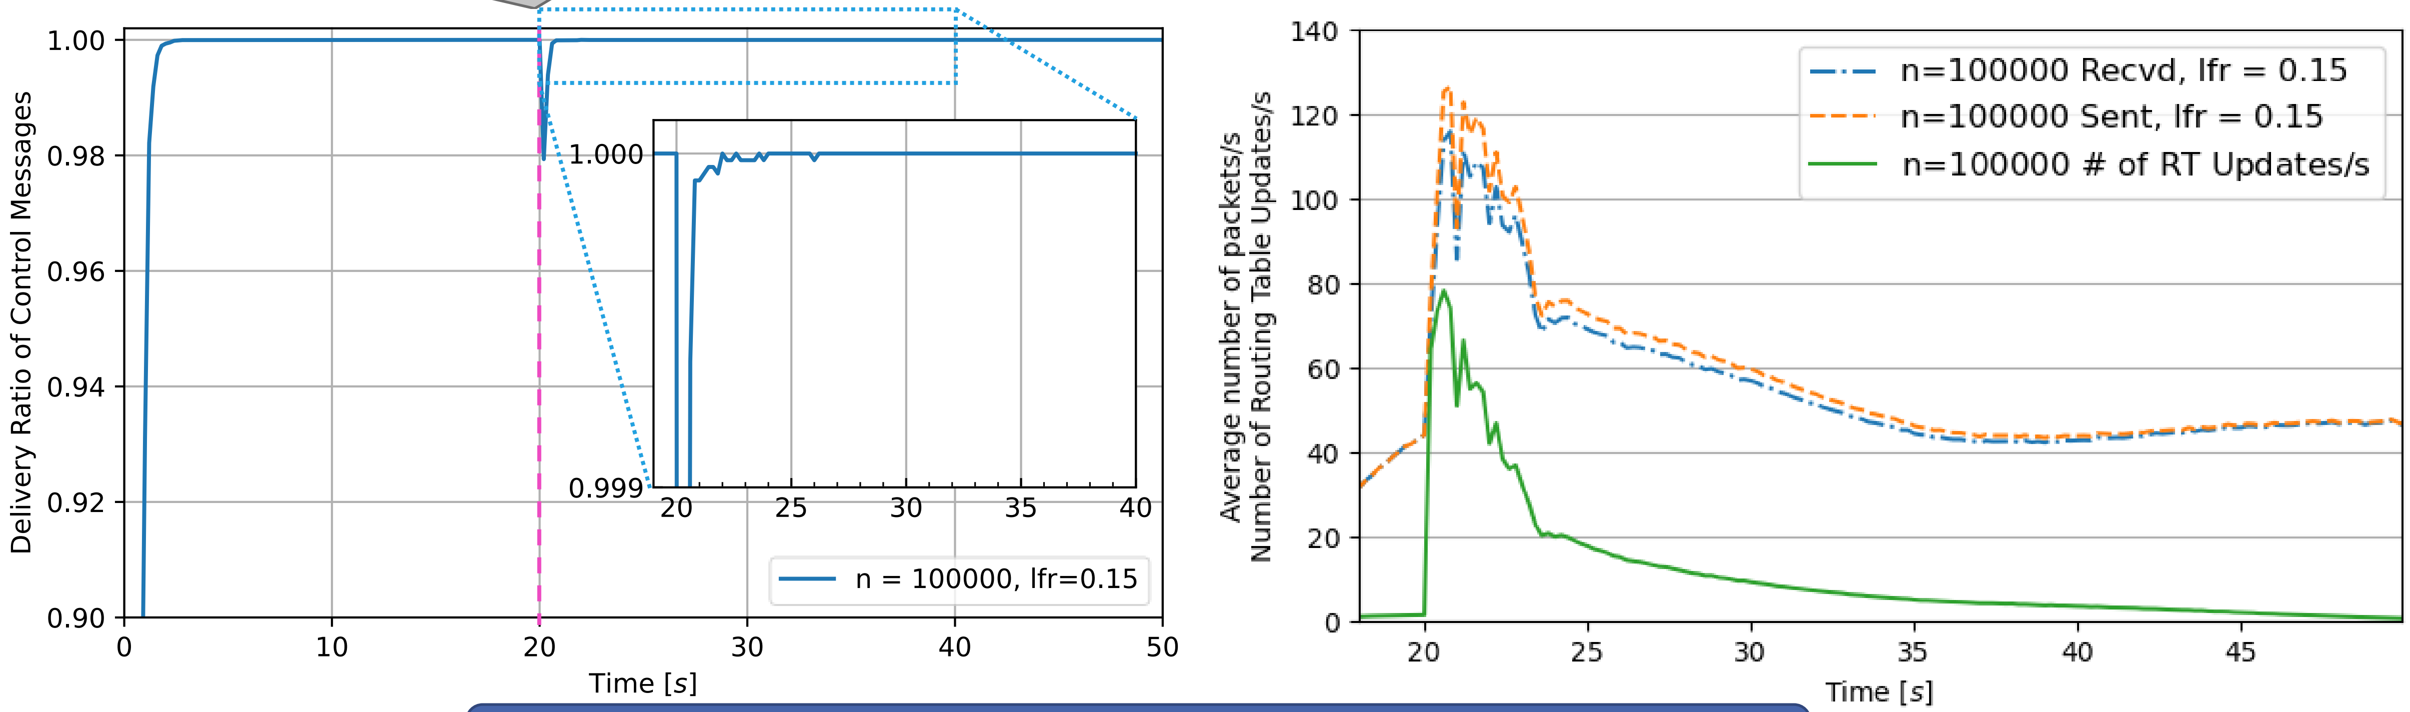


**Fast Convergence + Scalable Overhead**

# Conclusions

- KIRA

  provides **self-organized zero-touch** control plane connectivity
  → foundation for autonomic and resilient networks

  - Not (yet) a replacement for OSPF/IS-IS/BGP
  - Designed for large provider domains (e.g., 5G, 6G) to work across multiple providers
  - Security design ongoing
- KIRA integrates a DHT for simple name resolution/service discovery
- Supports scalable and efficient topology discovery (→ KeLLy)
- Special end-systems mode → reduces overhead even more
- Supports multi-path routing and forwarding

https://s.kit.edu/KIRA