Achieving Convergence-Free Routing using Failure-Carrying Packets

Karthik Lakshminarayanan, Matthew Caesar, Murali Rangan,
Tom Anderson, Scott Shenker, Ion Stoica

UC Berkeley and University of Washington

Dealing with link failures

- Increasingly stringent demands on routing
  - Gaming, media, VoIP
- Problem: traditional routing paradigms deal poorly with link failures
  - Delayed convergence
  - Update storms
  - Externally-visible churn

Many solutions exist

- Backup paths, reducing timers, speeding up computation
- These reduce, but do not eliminate the convergence process
- **Is it possible to completely <u>eliminate</u> convergence from routing?**

Eliminating convergence with FCP

- Our approach: Failure-carrying packets
- Our focus: enterprise networks/ISPs
- Main idea:
  - Routers maintain map of links in network
  - Transient changes to this map (failures) are appended to data packets

Roadmap

- What is FCP?
- How does it work?
- How well does it perform?

Roadmap

- What is FCP?
- How does it work?
- How well does it perform?

Example: FCP routing

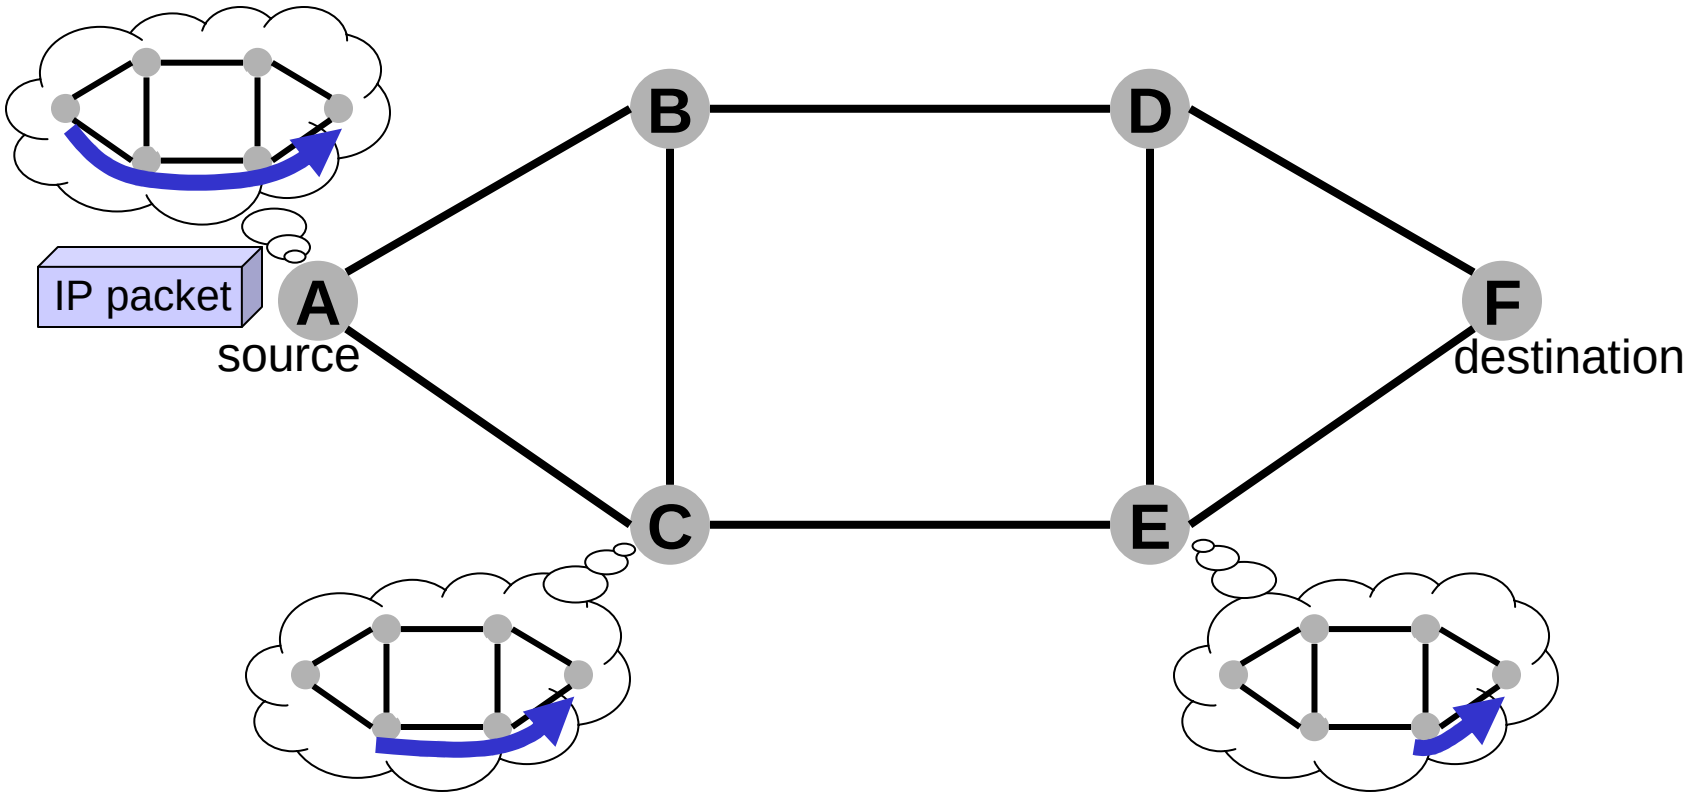

Example: FCP routing

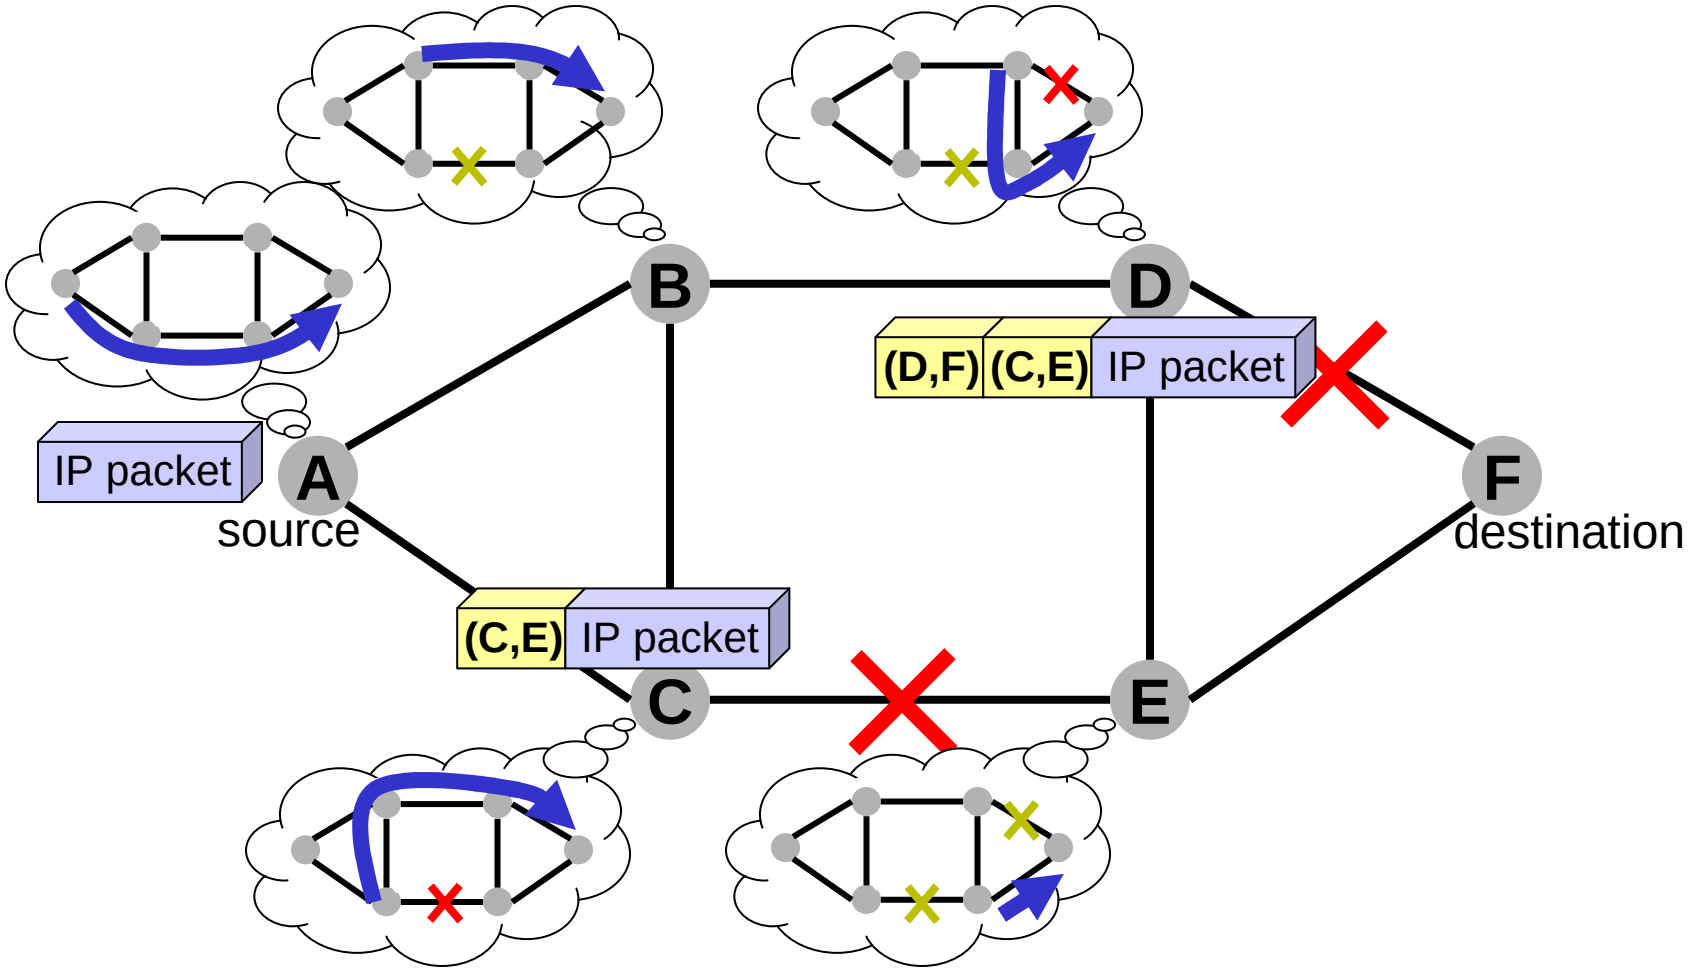

Properties of FCP

- Eliminates the convergence process, guarantees packet delivery
- Major conceptual change
- Quantitative performance improvement

Example: FCP doesn't find a path

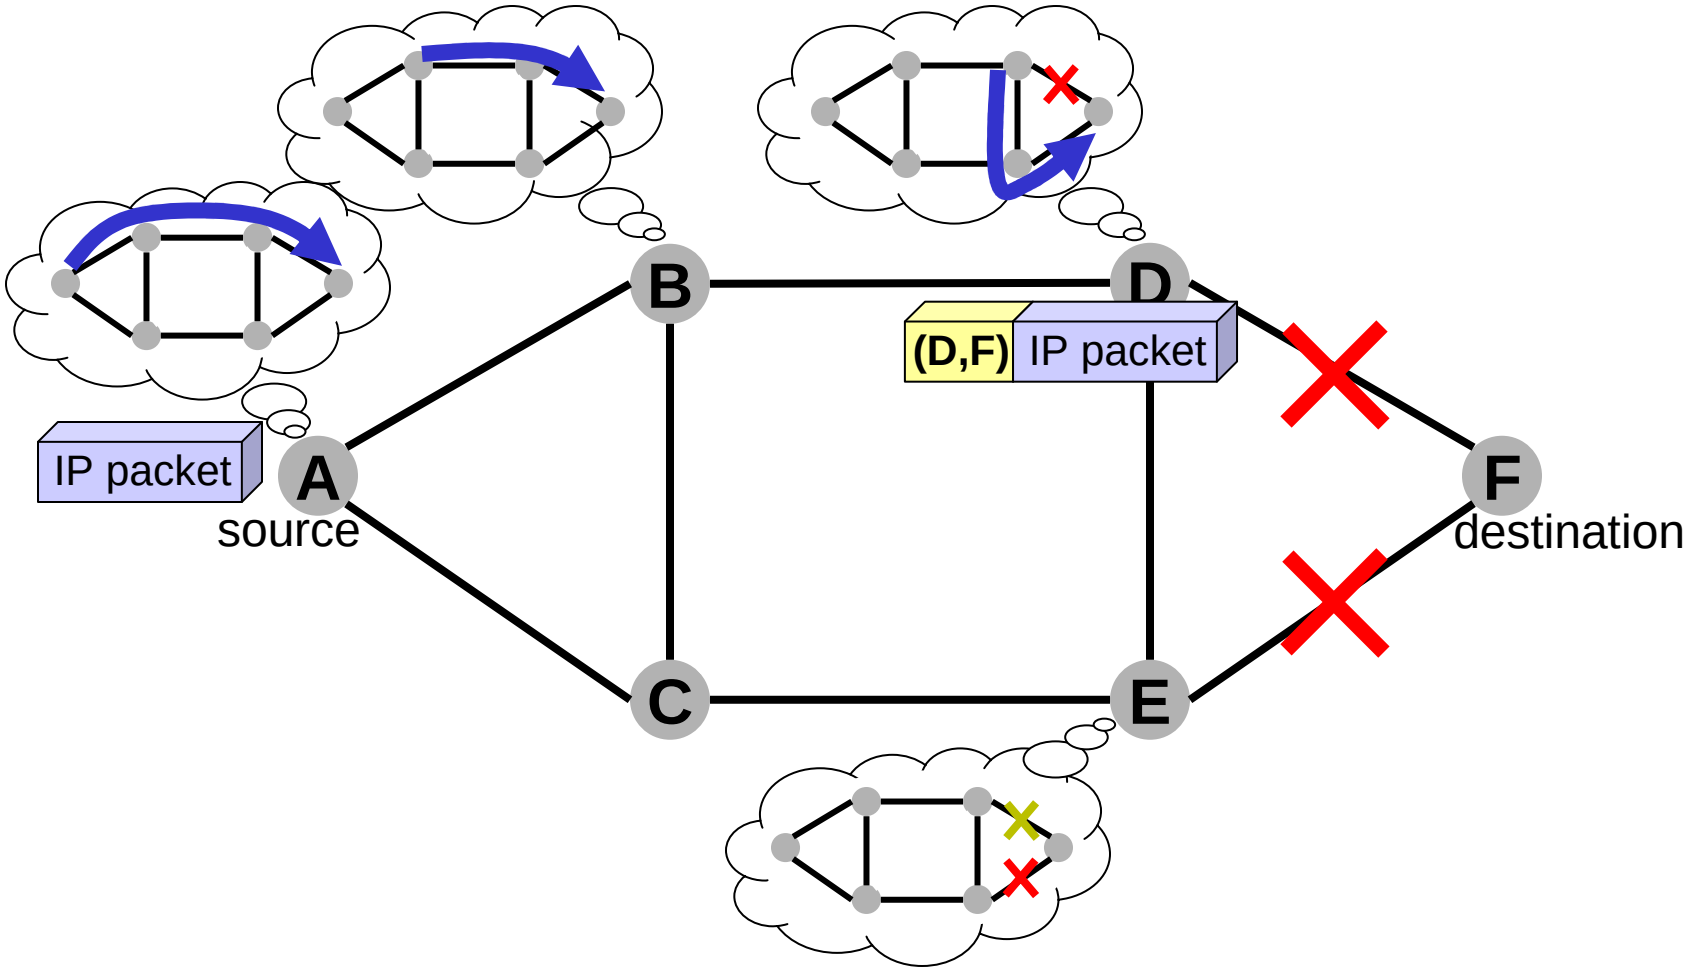

FCP guarantees reachability

- Define Liveness Graph LG(t1,t2) as set of links alive during entire time interval [t1,t2]
- Link propagation delay is 1 time unit
- No loss due to congestion or failure detection

- If the number of failures is less than
  (t2-t1)/diameter(LG)
  and if LG spans all nodes:
  Then a packet sent at t1 will arrive at time t2

Roadmap

- What is FCP?
- How does it work?
- How well does it perform?

Map dissemination

- Map of links distributed by a replicated coordinator, through reliable flooding
- Each map has sequence number, routers write this in packet headers during map transitions
- Alternative: eliminate coordinator by appending source-routes (SR-FCP)

FCP's other properties

- FCP improves security:

  A compromised router cannot impact a packet's route, unless it is already on that route

- FCP is resilient to inconsistency:

  If maps share a connected Liveness Graph, then a sent packet will arrive at its destination

Reducing overhead of FCP

- Several optimizations reduce performance penalties:

Cost of FCP	Optimization
Computing alternate routes	Precompute alternate routes
Route computation time	Incremental recomputation
Space in packet header	Replace failed links in header with labels

Roadmap

- What is FCP?
- How does it work?
- How well does it perform?

Evaluation setup

- Compared with OSPFD software router, Juniper's backup-path selection algorithm
- Configuration:
  - OSPF: hello_interval=400ms, dead_interval=2sec
  - FCP, Backup-paths: hello_interval=50ms, dead_interval=250ms
- Replayed Internet2 IS-IS traces, synthetic updates on Rocketfuel topologies

Results: costs of FCP

- Packet header overhead:
  Avg. 0.03 bytes, max 8 bytes per packet

- Recomputation time:
  2% of pkts require on-demand recomputation
  Requires 21μs on average, 810μs max

- Stretch penalty:
  Average stretch of 1.07, max of 1.75

Results: benefits of FCP

- Lower loss:
  Increases availability during failure events from one 9 to three 9's
- Less control overhead:
  No control updates for transient failures
- Reduced state requirements:
  one to three orders of magnitude less state than backup-paths

Results: OSPF vs. FCP

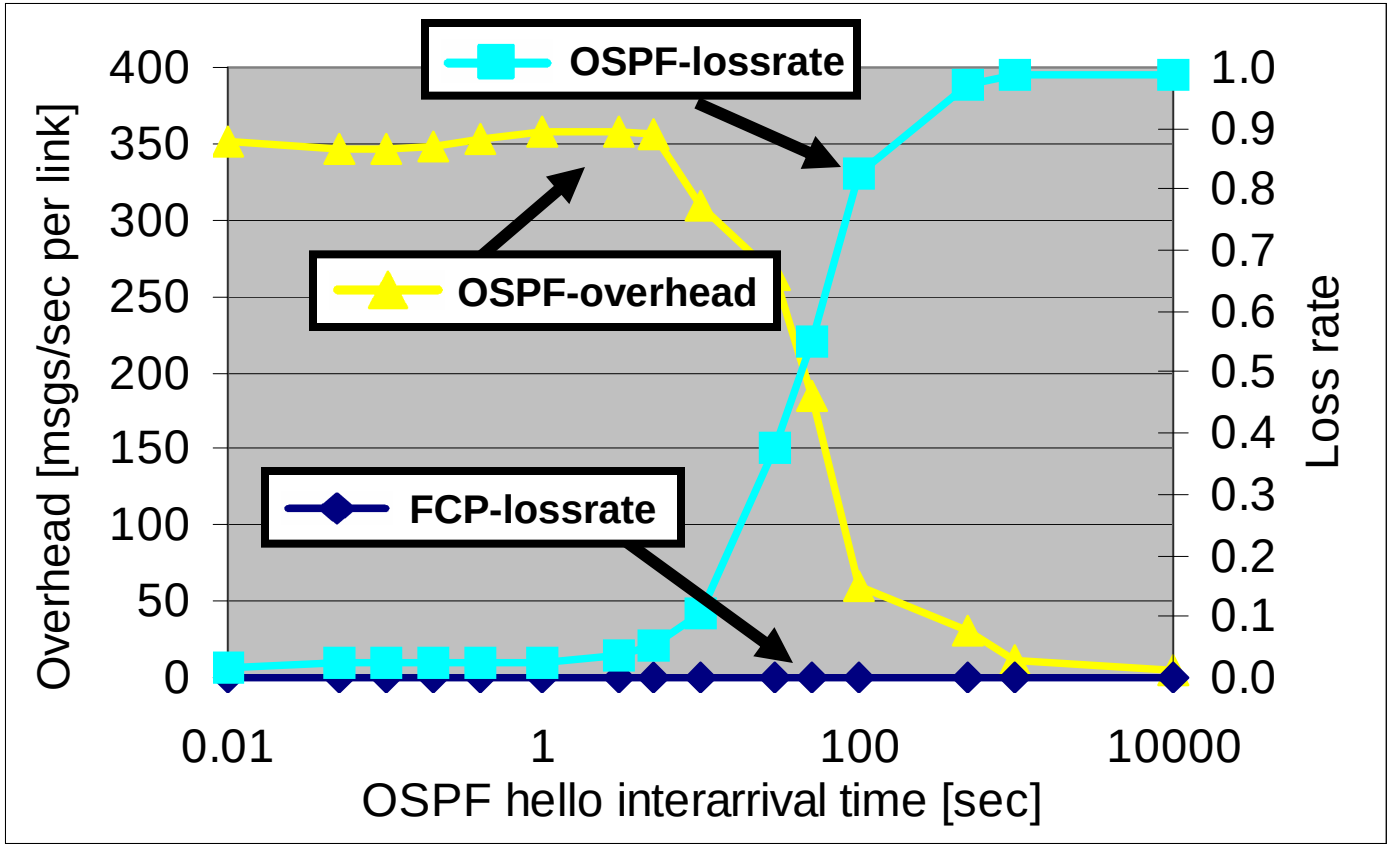

- Unlike FCP, OSPF cannot simultaneously provide low churn and high availability

Results: Backup-paths vs. FCP

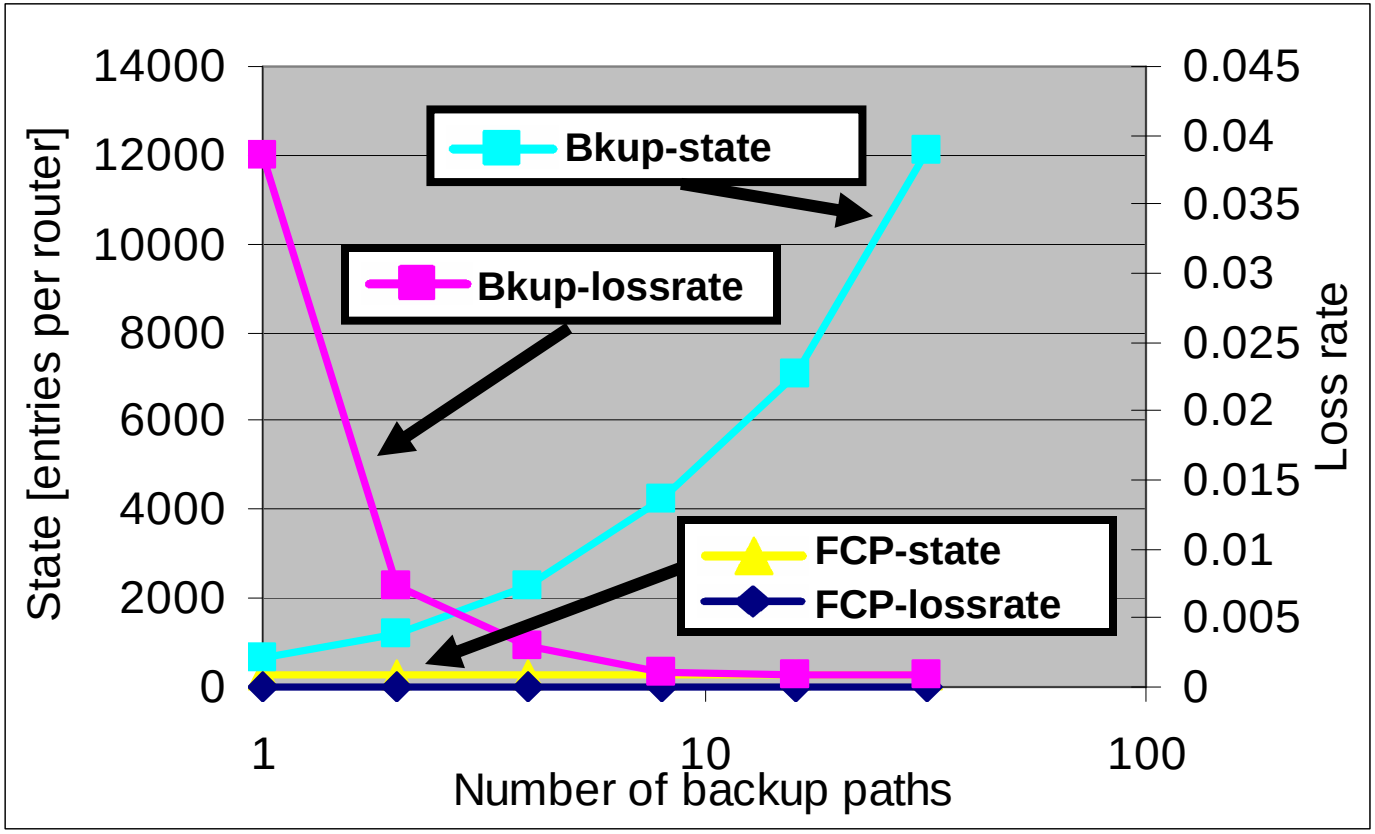

- Unlike FCP, Backup-paths cannot simultaneously provide low state and lossrate

Other applications of FCP

- Improving iBGP stability
  - Treat egress routes to prefixes as virtual links
  - On virtual link failure, forward to alternate egress
- Interdomain BGP routing
  - Use SR-FCP instead of centralized coordinator
  - Treat policy violations as link failures

Conclusions

- FCP represents a significant conceptual change
  - No need for instantaneous state consistency between routers
- FCP has several performance benefits
  - Lowers loss, control overhead, state
  - Stretch, computation costs are small
- Future work
  - FCP for sensornets, overlay networks