One Tunnel is (Often) Enough

Simon Peter, Umar Javed, Qiao Zhang, Doug Woos,
Thomas Anderson, Arvind Krishnamurthy

University of Washington

Financial support: NSF, Cisco, and Google

Internet Routing Has Issues

Outages and poor performance due to:

- Pathological routing policies
- Route convergence delays
- Misconfigured routers
- Prefix hijacking
- Malicious route injection (route table overload)
- Distributed denial of service

Good technical solutions in most/all cases, but glacial progress towards adoption.

Why?

The fault lies not in our stars, but in ourselves.

– Cassius

Why?

The wheels of justice grind slow but grind fine.

– Sun Tzu

Why?

We don’t care, we don’t have to, we’re the phone company.

– Lily Tomlin

Local Problem => Global Outage

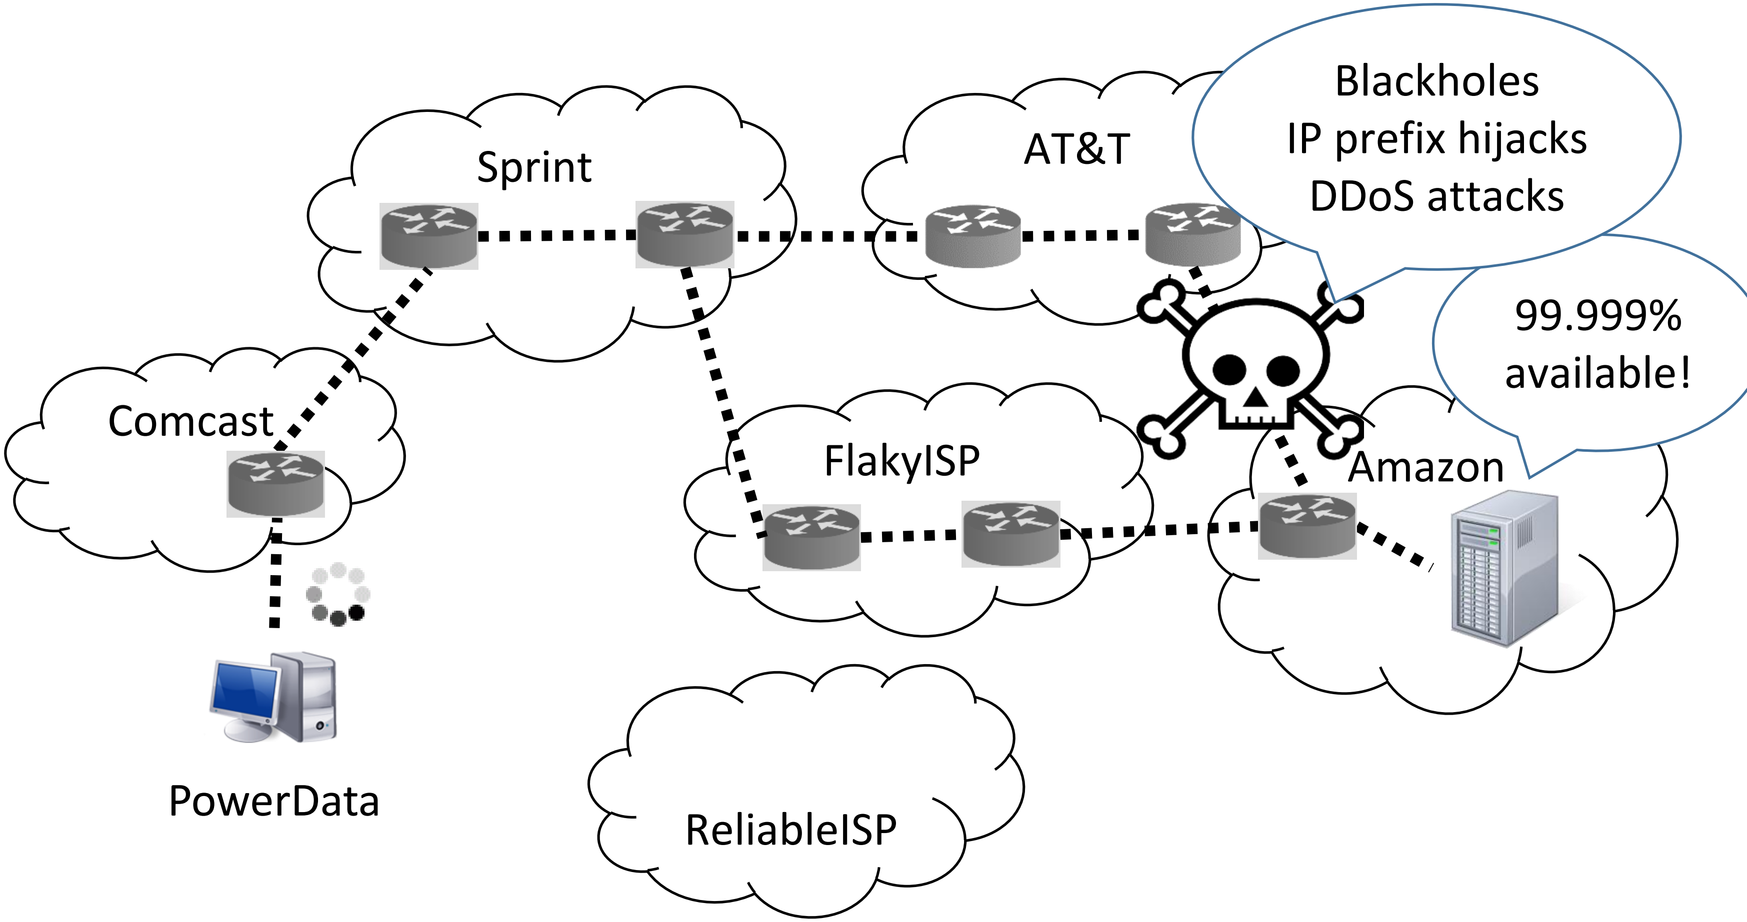

Local Problem => Global Outage

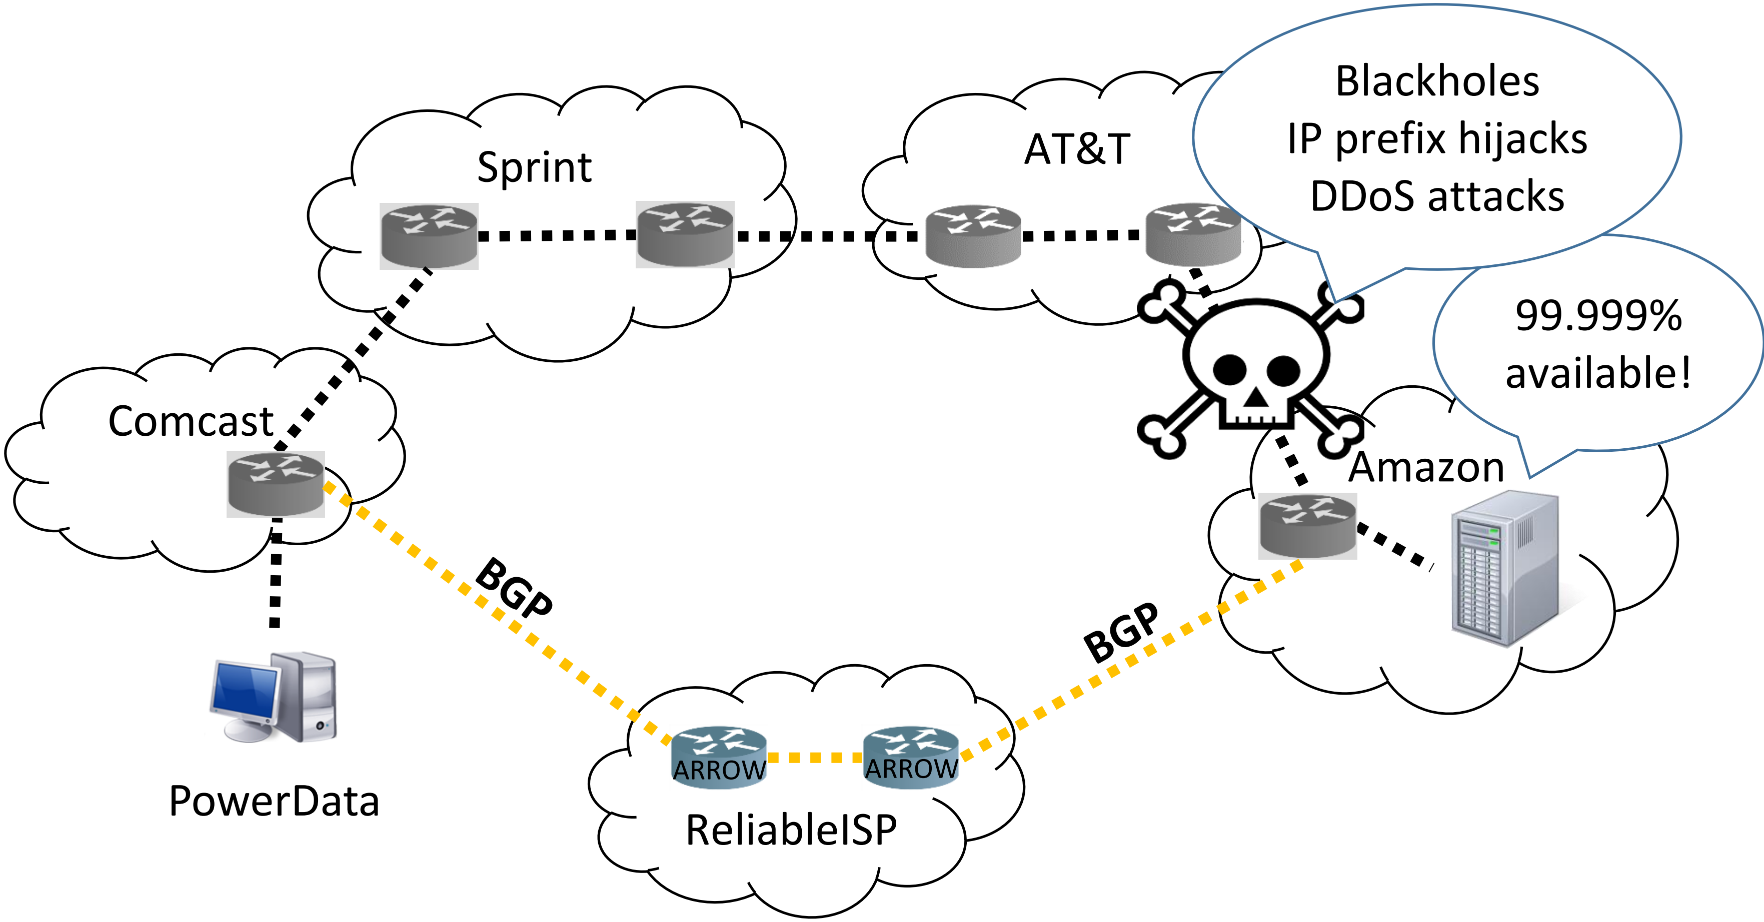

Can we turn local reliability into global reliability?

Assumptions/Observations

Shorter paths are more reliable than longer paths

Simple packet processing is feasible at high-speed border routers

- 10 Gb *per core* on commodity hardware

AS graph is relatively small and stable

See paper for quantitative justifications.

ARROW

- ARROW: Advertised Reliable Routing Over Waypoints
  - ISPs offer a QoS tunnel across their network to remote customers
  - Paid service akin to AWS or Google Cloud
- ARROW runs on a small ISP we control
- Evaluation: ARROW effective even if only a single tier-1 ISP adopts

ARROW Example

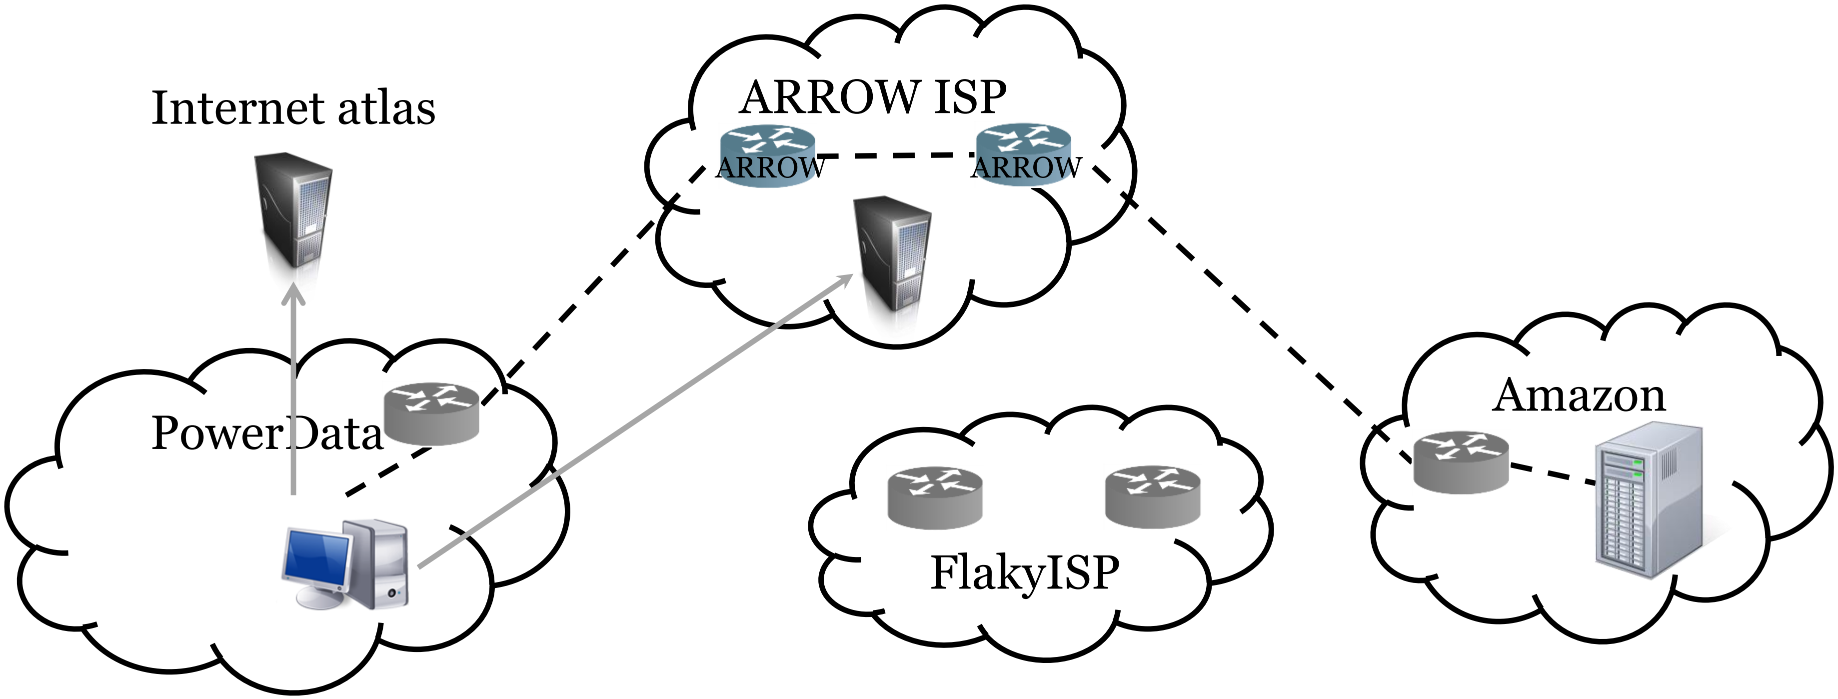

1. Consult atlas of ISPs offering ARROW services
2. Construct tunnel through ARROW ISP, to output target address

Use Cases

- Enterprises
  - More reliable access to cloud services
  - QoS between physically remote locations
  - Home health monitoring
- Business-facing ISP or cellular telecom
  - Market share driven by perceived data network performance, reliability
  - Well-developed market for premium service
  - 70% of data traffic exits telecom network

ARROW Mechanisms

How does endpoint/proxy know what tunnels are available?

- Atlas published by ISPs offering ARROW service: latency/bw/cost to which prefixes
- User/app-specific path selection

How are packets encapsulated?

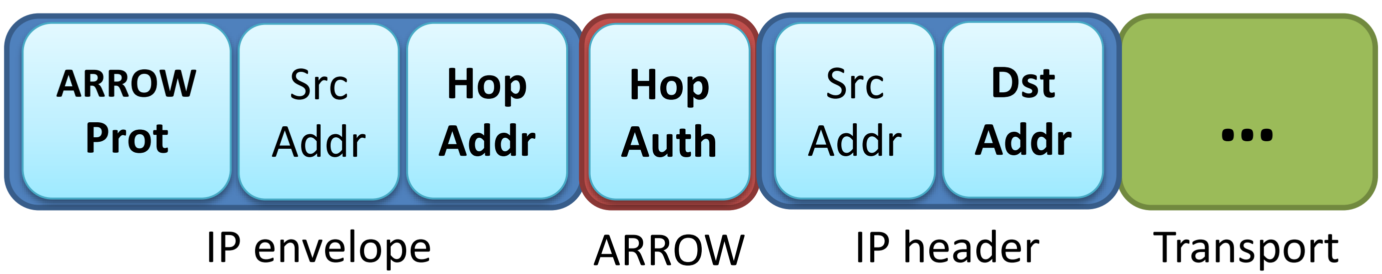

ARROW Mechanisms

How are packets authenticated?

- Packet authenticator provided by ISP at setup
- Authenticator can be hashed with checksum of packet to prevent snoop-stealing

What ISP data plane operations are needed?

- Check authenticator
- Check packet is within rate limit envelope
- Handle fault isolation probe, if any
- Re-write destination address

Failover

- Failure of a router internal to an ISP
  - ARROW is a stateful service
  - Local detection/recovery using Zookeeper
- Failure of a border router
  - End system/proxy detection/recovery
  - Use backup route through another PoP
- Failure of an entire ISP
  - End system/proxy detection/recovery
  - Use backup route through other ISPs

Failure Isolation

How does endpoint/proxy locate who is at fault for service disruption?

- Send probe packet to locate the failure
- Each hop:
  - Responds to the previous hop
  - Forwards the probe packet to next hop

See UW TR for efficient Byzantine-resilient solution

Implementation

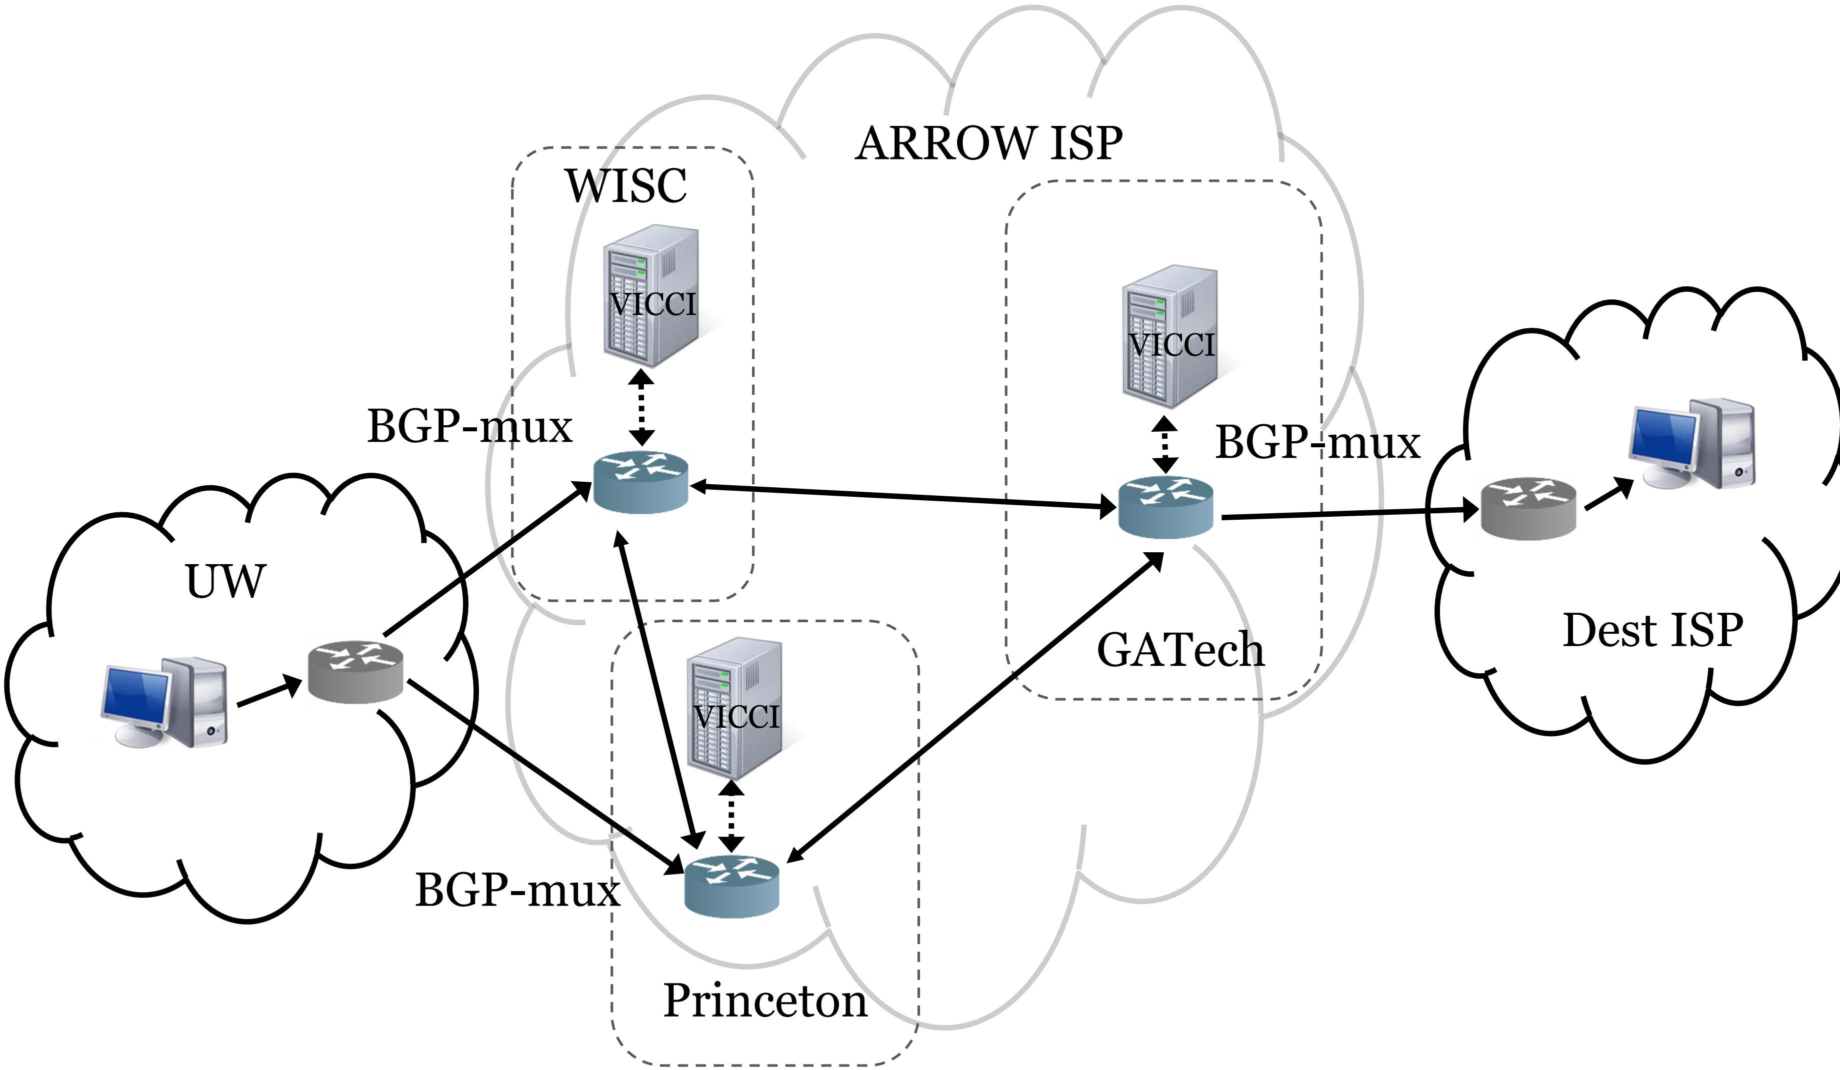

Overhead

What is the data plane overhead of ARROW?

	RTT (us)	Throughput (Gbps)
UDP/TCP	96	9.4
Serval	81	9.5
ARROW 1 hop	132	9.5

Failover Latency

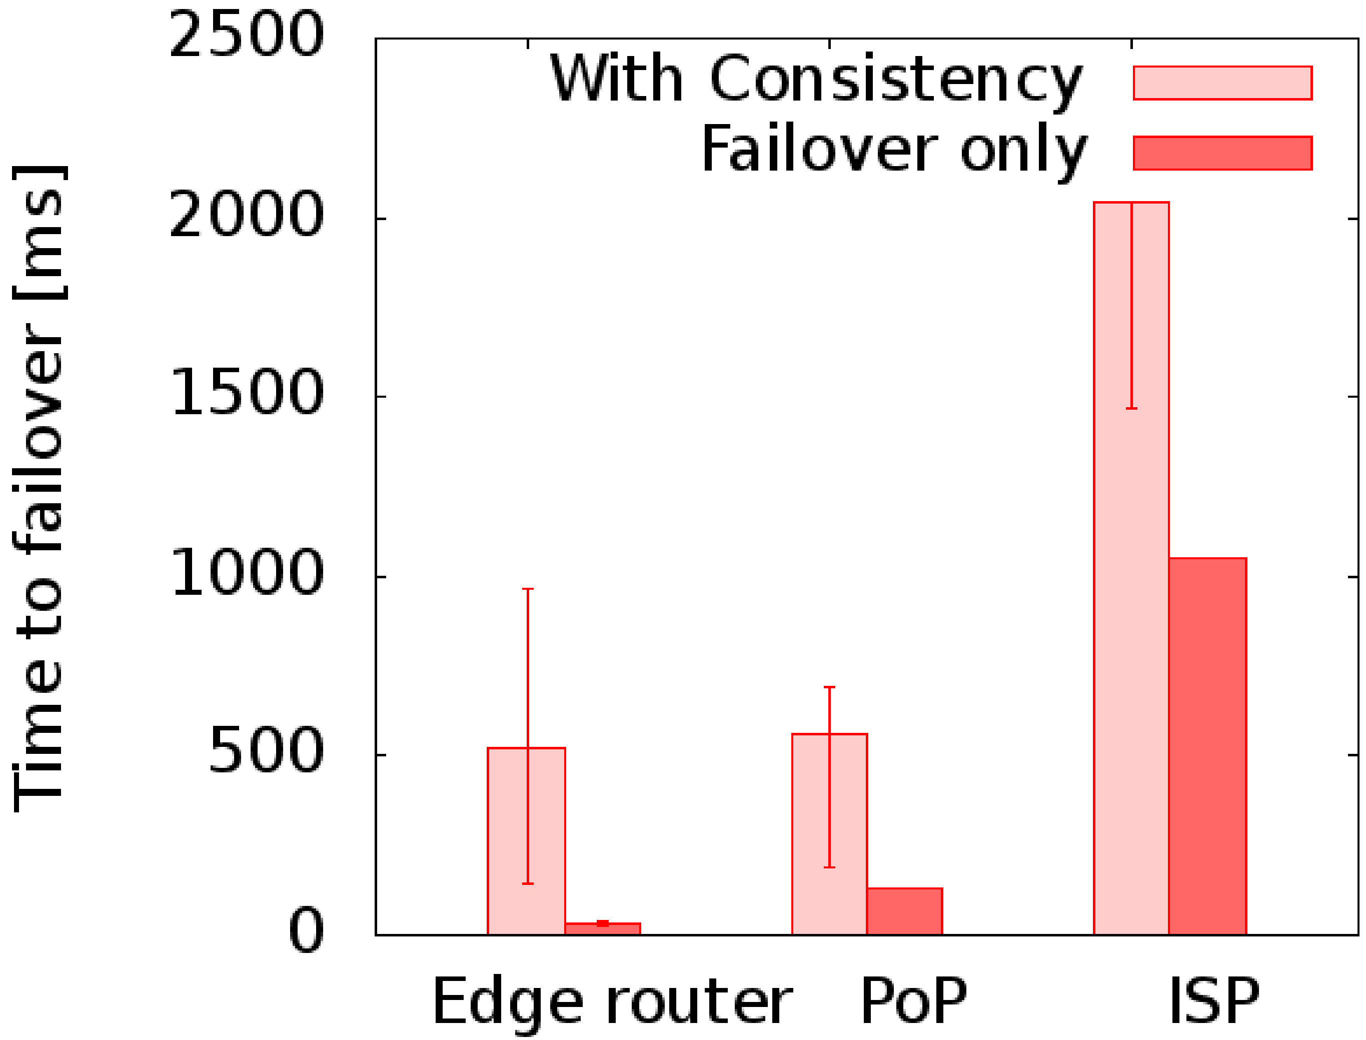

BGP Outage

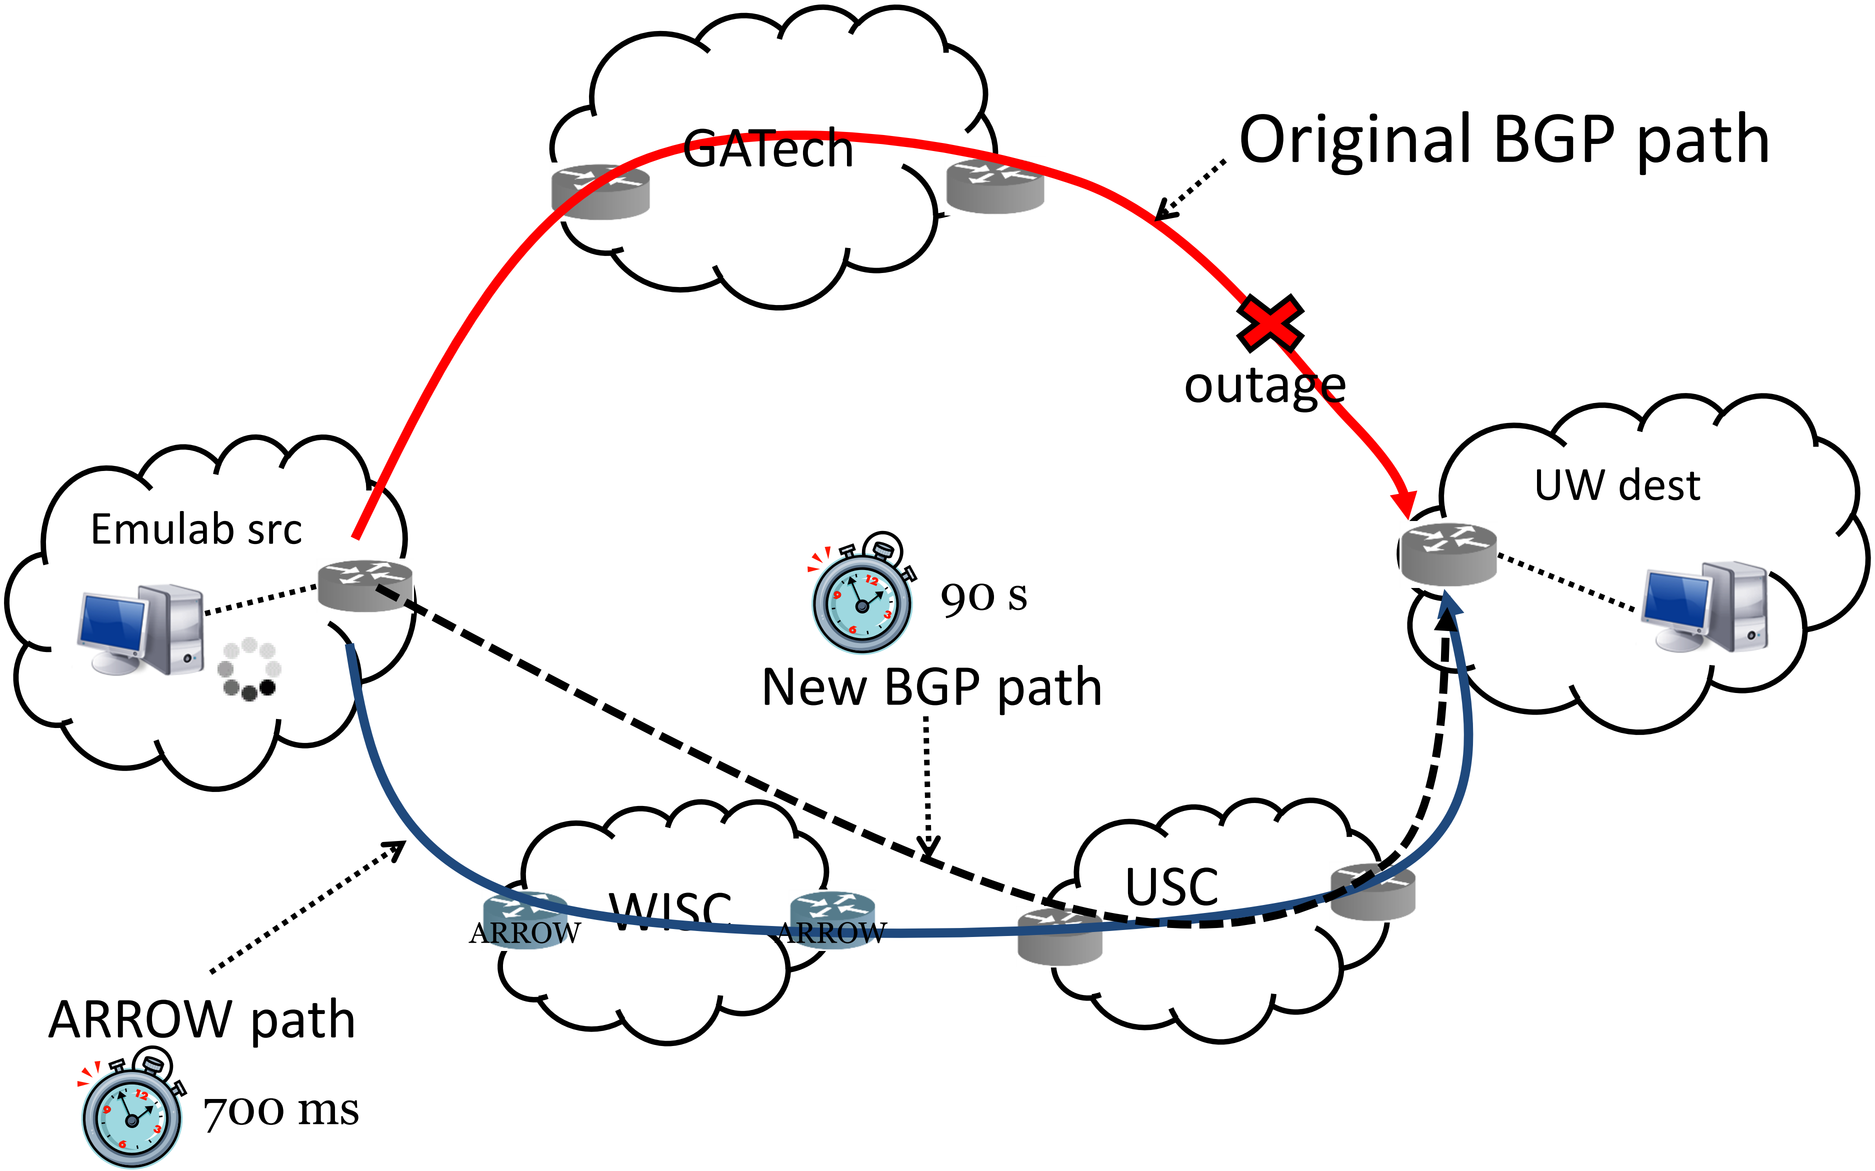

GATech
Original BGP path
outage
Emulab src
UW dest
90 s
New BGP path
ARROW
WISC
ARROW
USC
ARROW path
700 ms

Link Failure Disconnections (Simulated: 1 ARROW Tier-1)

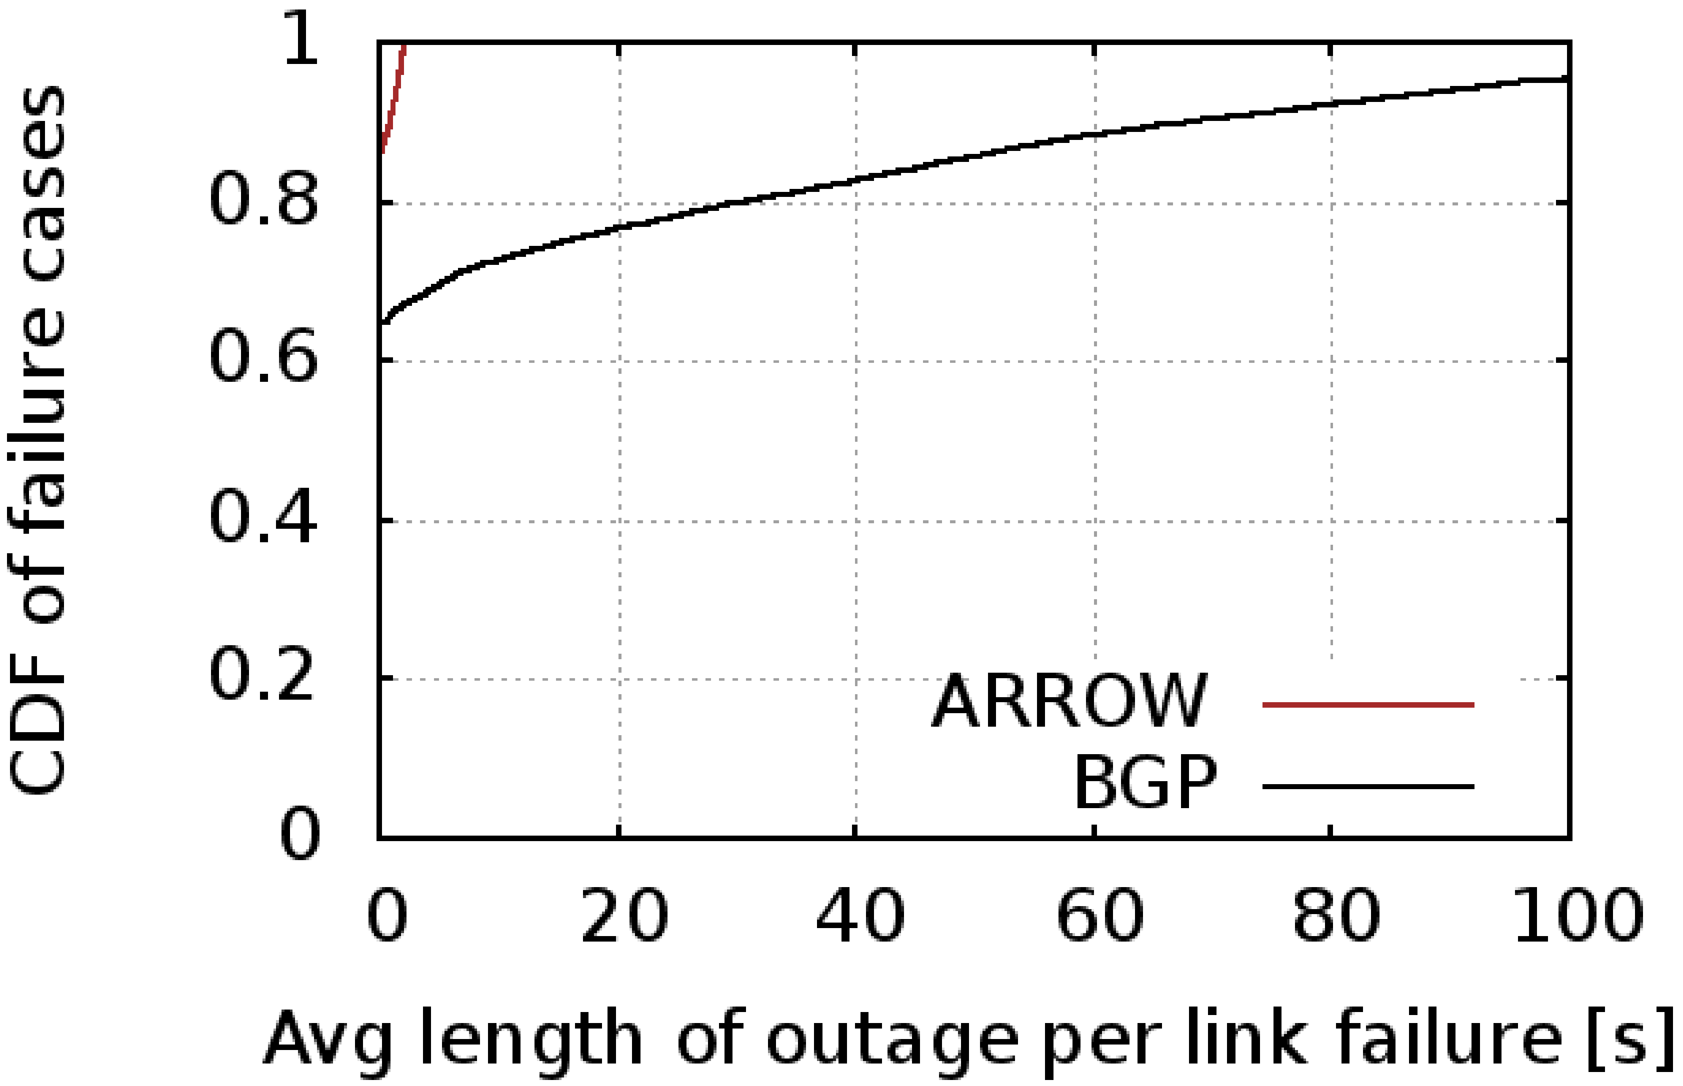

Prefix Hijacking (Simulated)

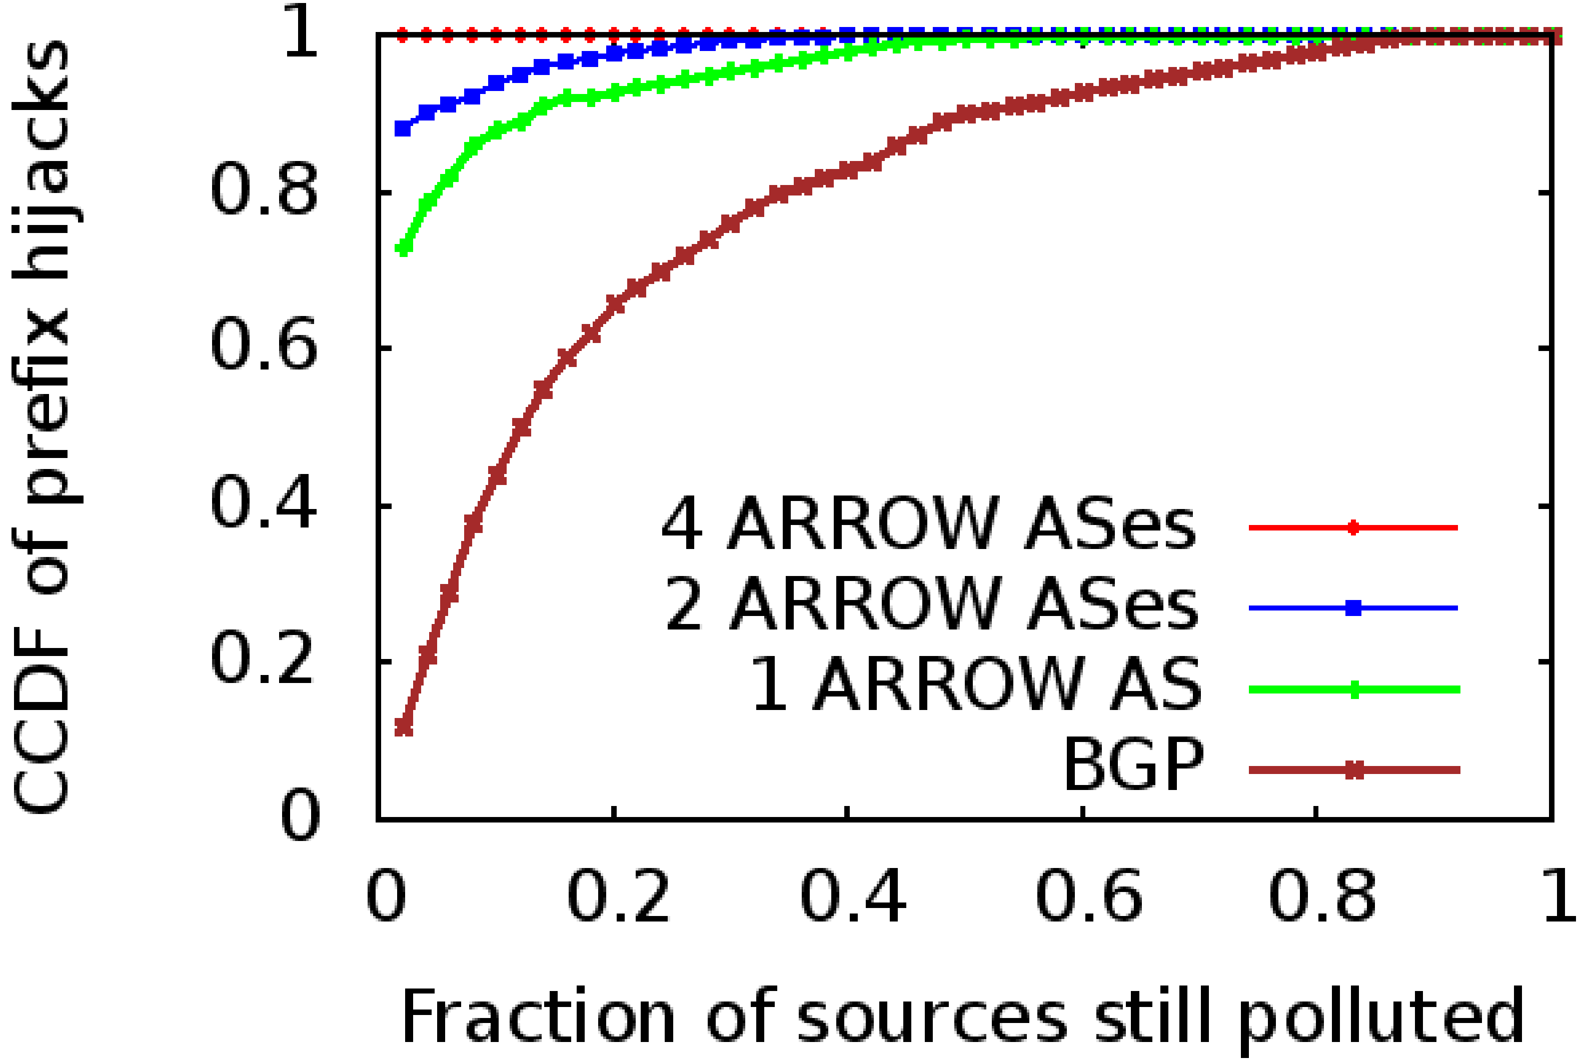

Summary

- ARROW: Advertised Reliable Routing Over Waypoints
  - ISPs offer a paid QoS tunnel across their network to remote customers
  - ARROW runs on a small ISP we control
  - Also on Google Cloud Platform
- One tunnel (through a tier-1) is often enough