Measuring and Inferring Weather's Effect on Residential Internet Infrastructure

Ramakrishna Padmanabhan, Aaron Schulman,
Ramakrishnan Sundara Raman, Reethika Ramesh,
Dave Levin, Neil Spring

Residential Internet infrastructure is vulnerable to weather

- Residential Internet infrastructure includes
  - power infrastructure
  - last-mile infrastructure from the ISP
- Lightning strikes, wind, and rain, can damage these infrastructure

Are there regions where the Internet infrastructure is particularly vulnerable?

- Some geographic regions may be particularly prone to weather-related Internet outages
  - E.g.: Snow in the southern U.S. states
- Identifying vulnerable regions and networks will:
  - help isolate underlying challenges
  - inform which enhancements can improve residential Internet reliability

Study residential Internet infrastructure reliability across regions using ThunderPing

- Internet reliability can be studied by detecting and analyzing outages
- ThunderPing detects residential Internet outages in the U.S.
  - across geography, ISPs, linktypes
  - in many weather conditions

ThunderPing detects outages during times of predicted severe weather in the U.S.

Weather Alert

<entry>

<title>Hurricane Warning issued September 10 at 3:30AM EDT until September 10 at 11:30AM EDT by NWS</title>
…
<cap:severity>Severe</cap:severity>
<cap:certainty>Possible</cap:certainty>
…
<valueName>FIPS6</valueName>
<value>12003</value>
…
</entry>

Identify locations to ping

Maxmind Geolocation

…
24.128.255.0/24: 42.2843,-85.2293
24.129.0.0/23: 30.2187,-81.7540
…
24.129.42.0/23: 30.2558,-82.1300
…
76.123.95.224/27: 26.9856,-82.0910
76.123.96.0/23: 30.3533,-81.4990
…
76.123.116.0/23: 30.1294,-81.7775
…

Reverse DNS

c-24-129-42-9.hsd1.fl.comcast.net.
c-24-129-42-16.hsd1.fl.comcast.net.
c-24-129-42-19.hsd1.fl.comcast.net.
…
c-76-123-116-89.hsd1.fl.comcast.net.
c-76-123-116-97.hsd1.fl.comcast.net.

Identify addresses to ping

PlanetLab

- Pl 0
- Pl 1
- Pl 2
- Pl 3
- Pl 4
- Pl 5
- Pl 6
- Pl 7
- Pl 8
- Pl 9

Ping from multiple vantage points

Probe Responses

24.129.42.9	S5 S0 S8 S1 S3 F9 F9 F9 F9 F9 F2 F4
24.129.42.16	S5 S0 S8 S1 S3 F9 F9 F9 F9 F9 F2 F4
24.129.42.19	S5 S0 S8 S1 S3 F9 F9 F9 F9 F9 F2 F4
24.129.42.20	S5 S0 S8 S1 S3 F9 F9 F9 F9 F9 F2 F4
24.129.42.21	S5 S0 S8 S1 S3 F9 F9 F9 F9 F9 F2 F4
24.129.42.22	S5 S0 S8 S1 S3 F9 F9 F9 F9 F9 F2 F4
24.129.42.47	S5 S0 S8 S1 S3 F9 F9 F9 F9 F9 F2 F4
76.123.113.127	S8 S0 S1 S3 S9 F2 F2 F4 F6 F4 F2 F6
76.123.113.126	S8 S0 S1 S3 S9 F2 F2 F4 F4 F6 F2 F6
76.123.113.153	S5 S8 S0 S1 S3 F9 S9 F2 F2 F4 F6 F4
76.123.113.212	S8 S0 S1 S3 S9 F2 F2 F4 F6 F4 F2 F6
76.123.113.249	S8 S0 S1 S3 S9 F2 F2 F4 F6 F4 F2 F6
76.123.116.32	S8 S0 S1 S3 S9 F2 F4 F2 F6 F4 F2 F0
76.123.116.34	S8 S0 S1 S3 S9 F2 F4 F2 F6 F4 F2 F0
76.123.116.44	S8 S0 S1 S3 S9 F2 F4 F2 F6 F4 F2 F0
76.123.116.63	S8 S0 S1 S3 S9 F2 F4 F2 F6 F4 F0 F2
76.123.116.89	S8 S0 S1 S3 S9 F2 F4 F2 F6 F4 F2 F0
76.123.116.97	S8 S0 S1 S3 S9 F2 F4 F2 F6 F4 F2 F0

If previously responsive address stops responding, infer outage

Associate weather with responsive times and with times when failures began

- Consider 10 addresses that were pinged for 24 hours
  - Clear weather: 18 hours
  - Rain: 6 hours
- Suppose each address had an hour-long failure in clear weather and in rain, but was responsive otherwise
- Responsive time in Clear: (18 - 1) * 10 = 170 hours
- Responsive time in Rain: (6 - 1) * 10 = 50 hours

We use failure rate as our metric for comparing reliability

- Failure rate: # Failures/Responsive-duration (in days)
  - Failure rate in clear: 10/(170/24) = 1.41
  - Failure rate in rain: 10/(50/24) = 4.8
- Normalized failure rate = 4.8/1.41 = 3.4
- Failure rate in rain is 3.4 times that of clear

Short paper in IMC'11 presented preliminary results using 3 months' data

- Failures in Thunderstorm 4 times as likely as failures in Clear
- Failure in Rain twice as likely
- Since then, we have pinged for 7 years in a variety of weather conditions
- I will focus upon measurements from 2017

ThunderPing pings addresses in clear weather a lot

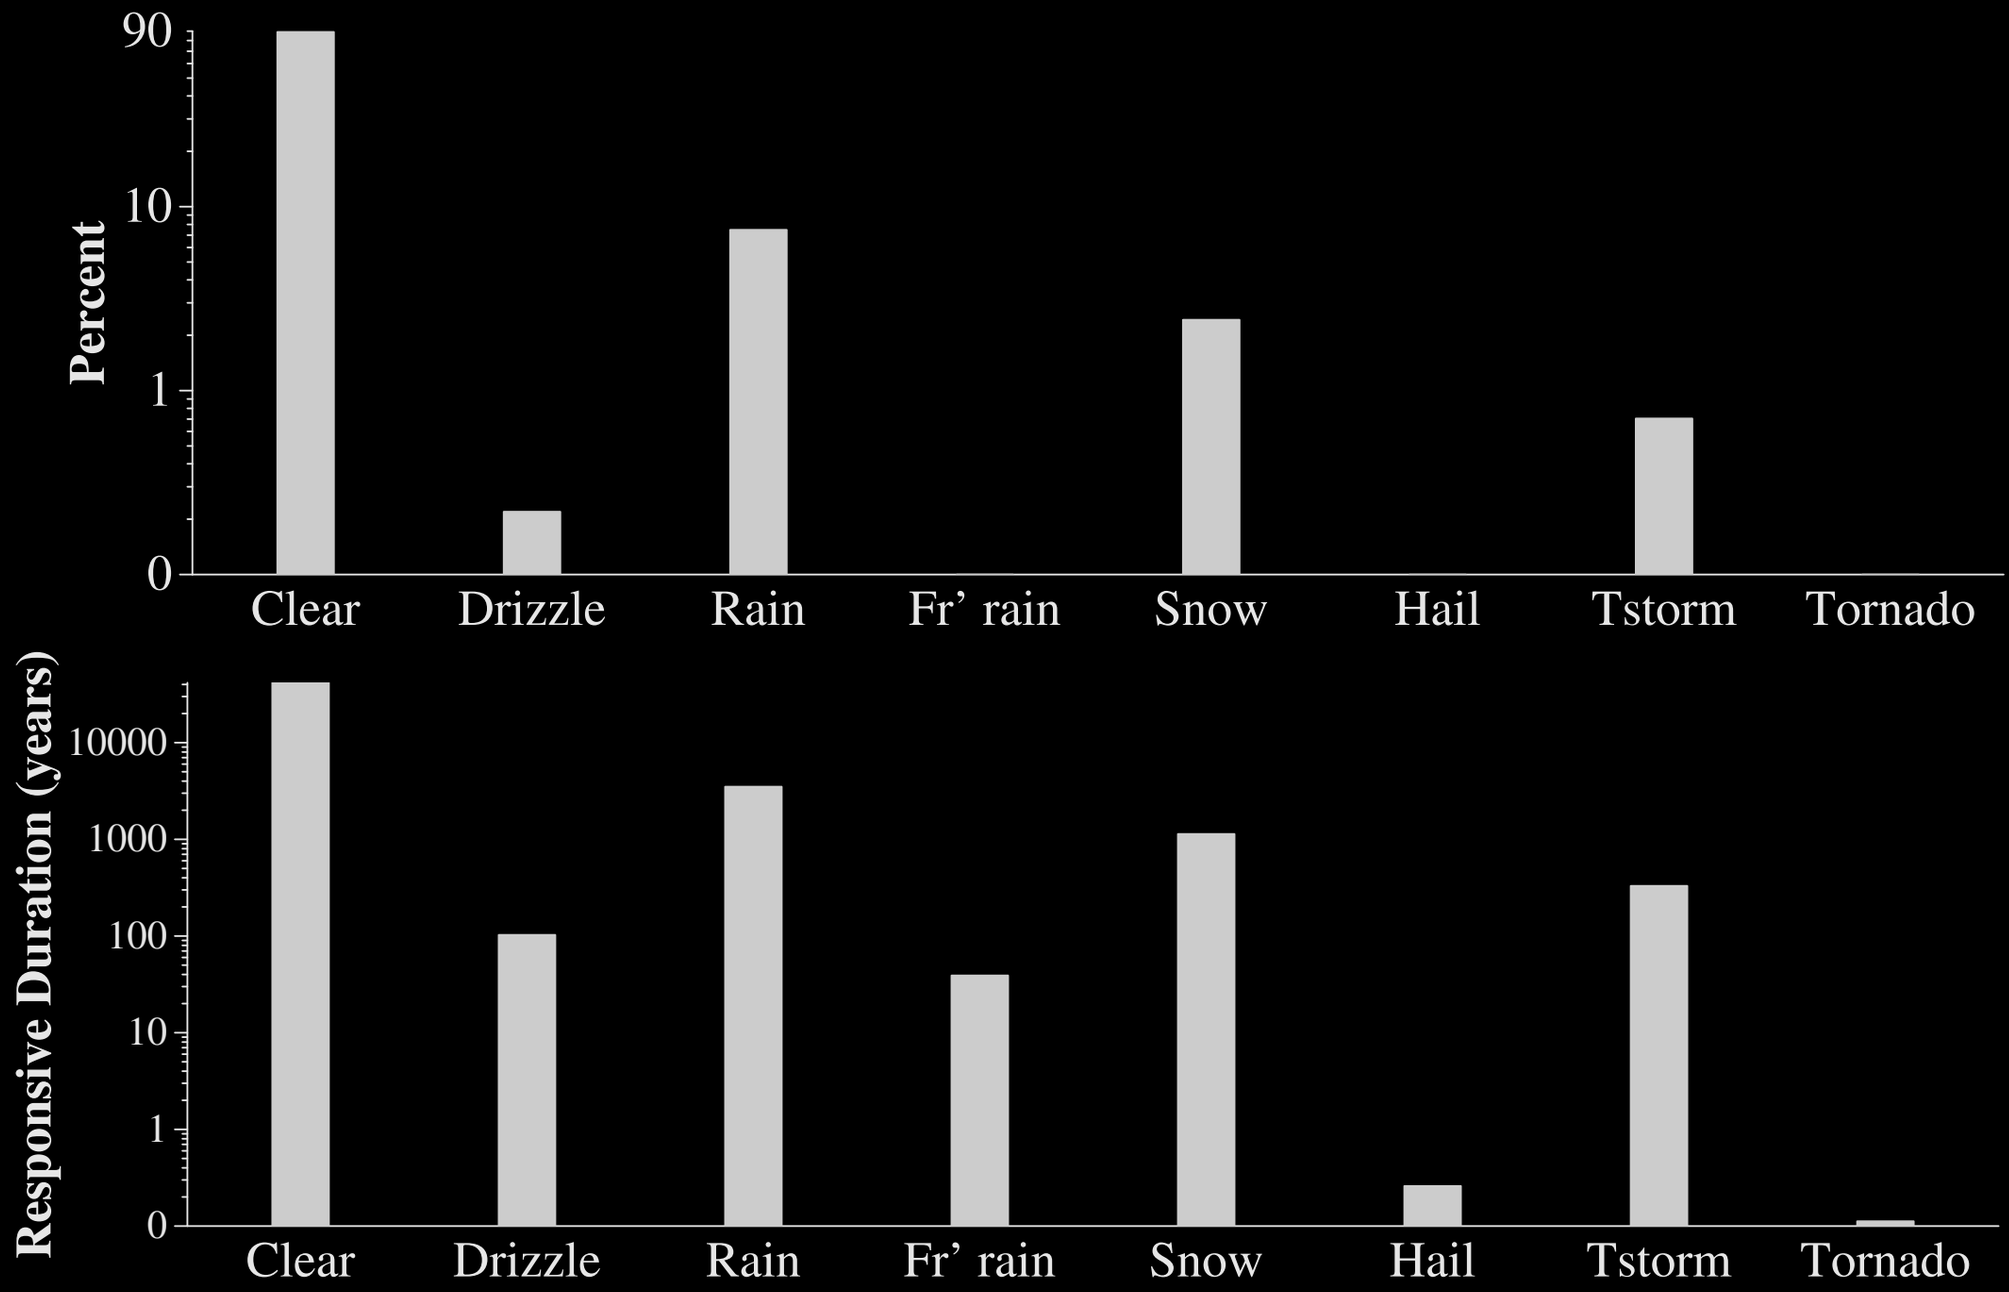

Distribution of addresses pinged in 2017

	ASN	AS Name	Pinged IPs	Failed IPs
Cable	Many	Comcast	1,118K	93K
	20115	Charter	288K	32K
	22773	Cox	62K	5K
DSL	209, 22561	Centurylink	476K	106K
	7029	Windstream	363K	55K
	701	Verizon (DSL)	104K	14K
Fiber	701	Verizon (Fiber)	173K	8K
Satellite	7155	Viasat	74K	47K
WISP	17306, 23205	RISE Broadband	22K	6K
		Rest	889K	147K
		Total	3,569K	513K

Failure rate in clear weather varies across linktypes

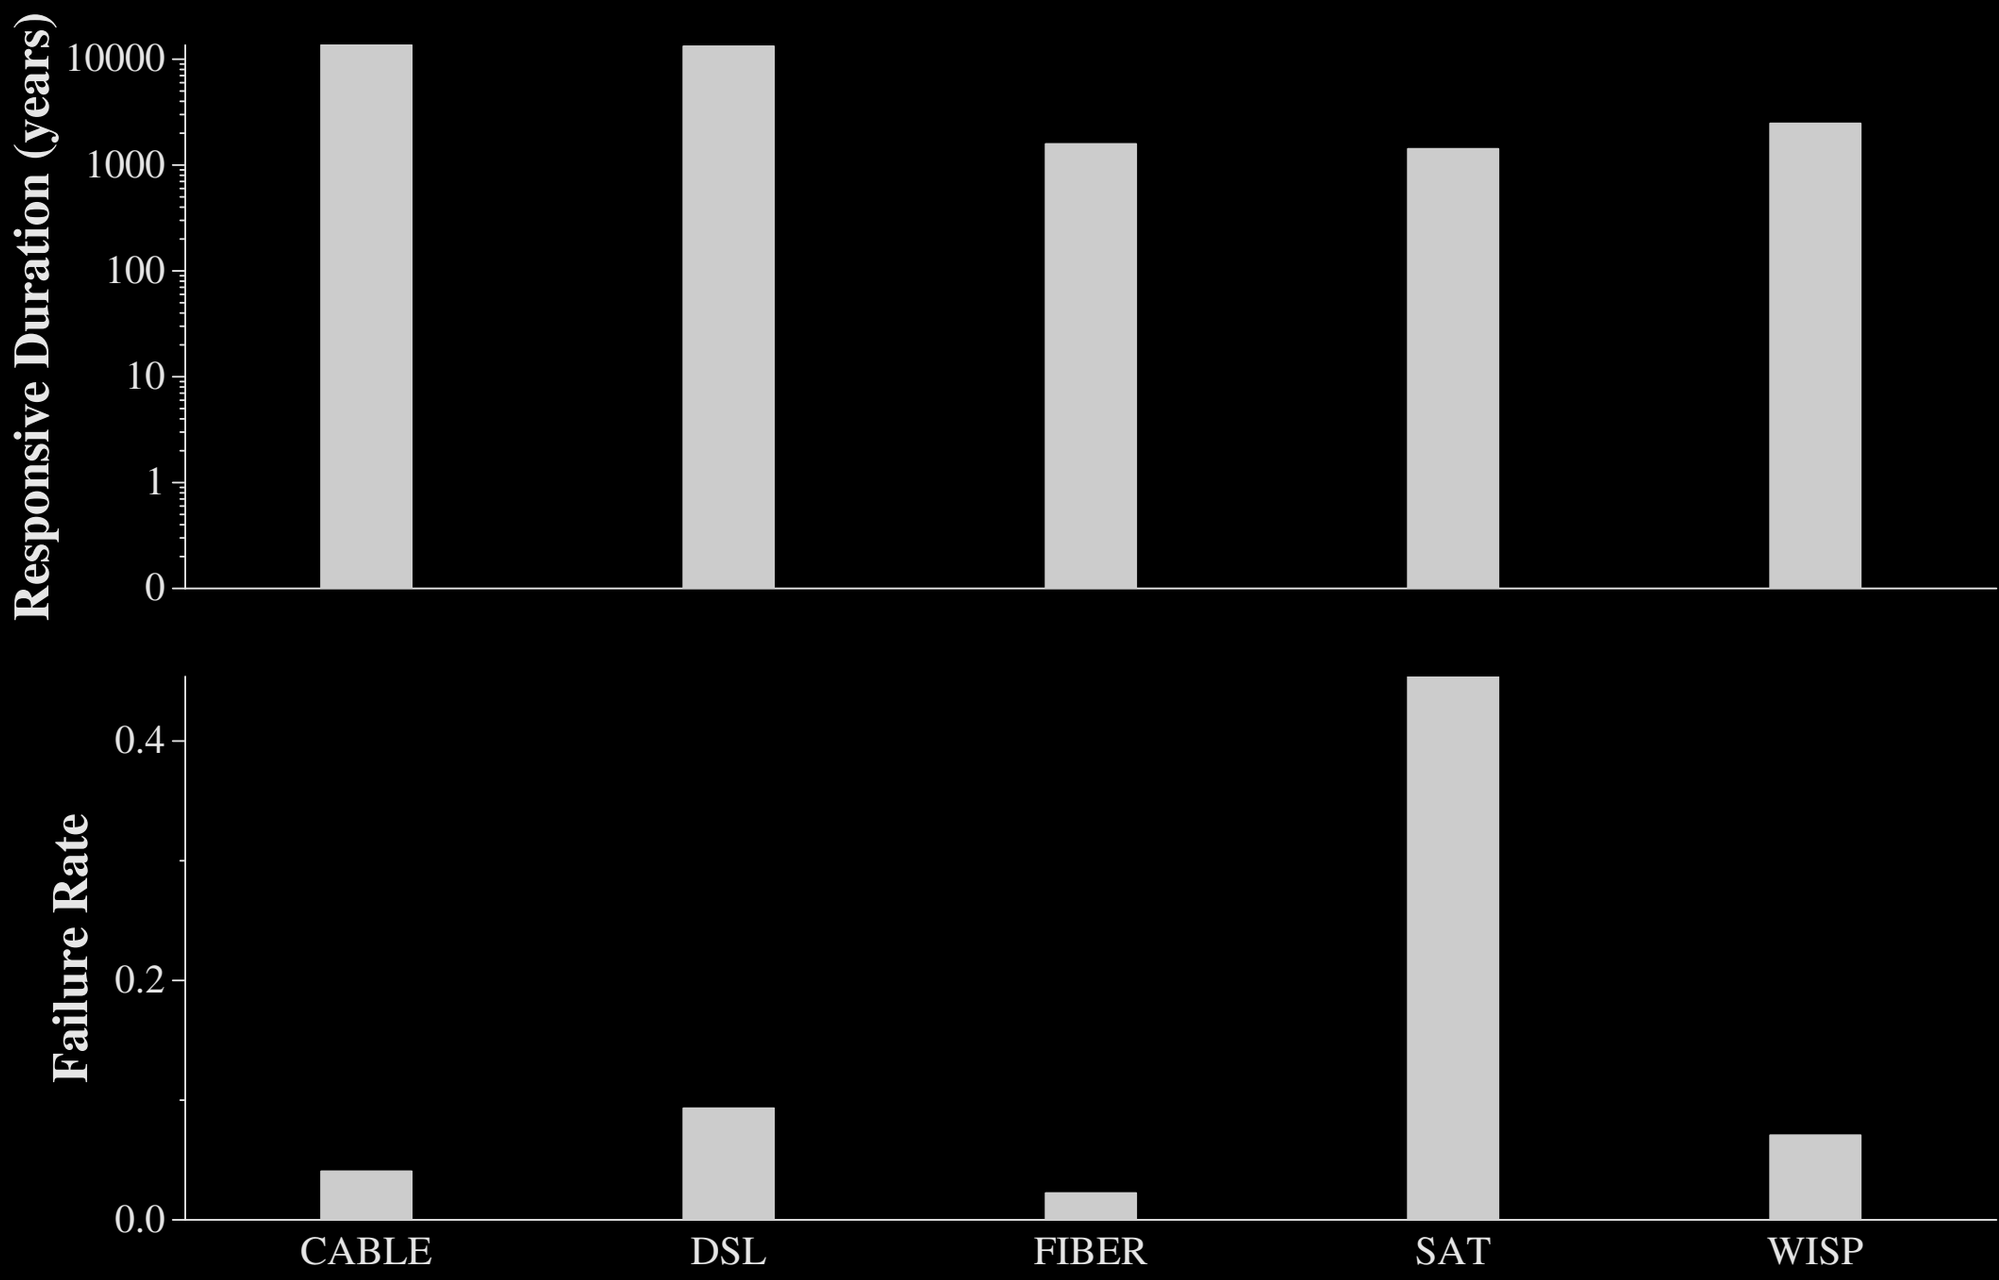

How does the failure rate in clear weather vary across regions?

- The failure rate for a linktype should be similar across regions
  - Modulo rural-urban divisions
- For the largest cable (Comcast) and DSL (Centurylink) ISPs in our dataset:
  - Found failure rates for all U.S. states
  - Plotted heat map of failure rates across U.S. states:
    - Heat map shows only states where addresses from the ISP were responsive for a total of 100 days (or more) in 2017

For Comcast, failure rate in clear weather is higher in the south-east

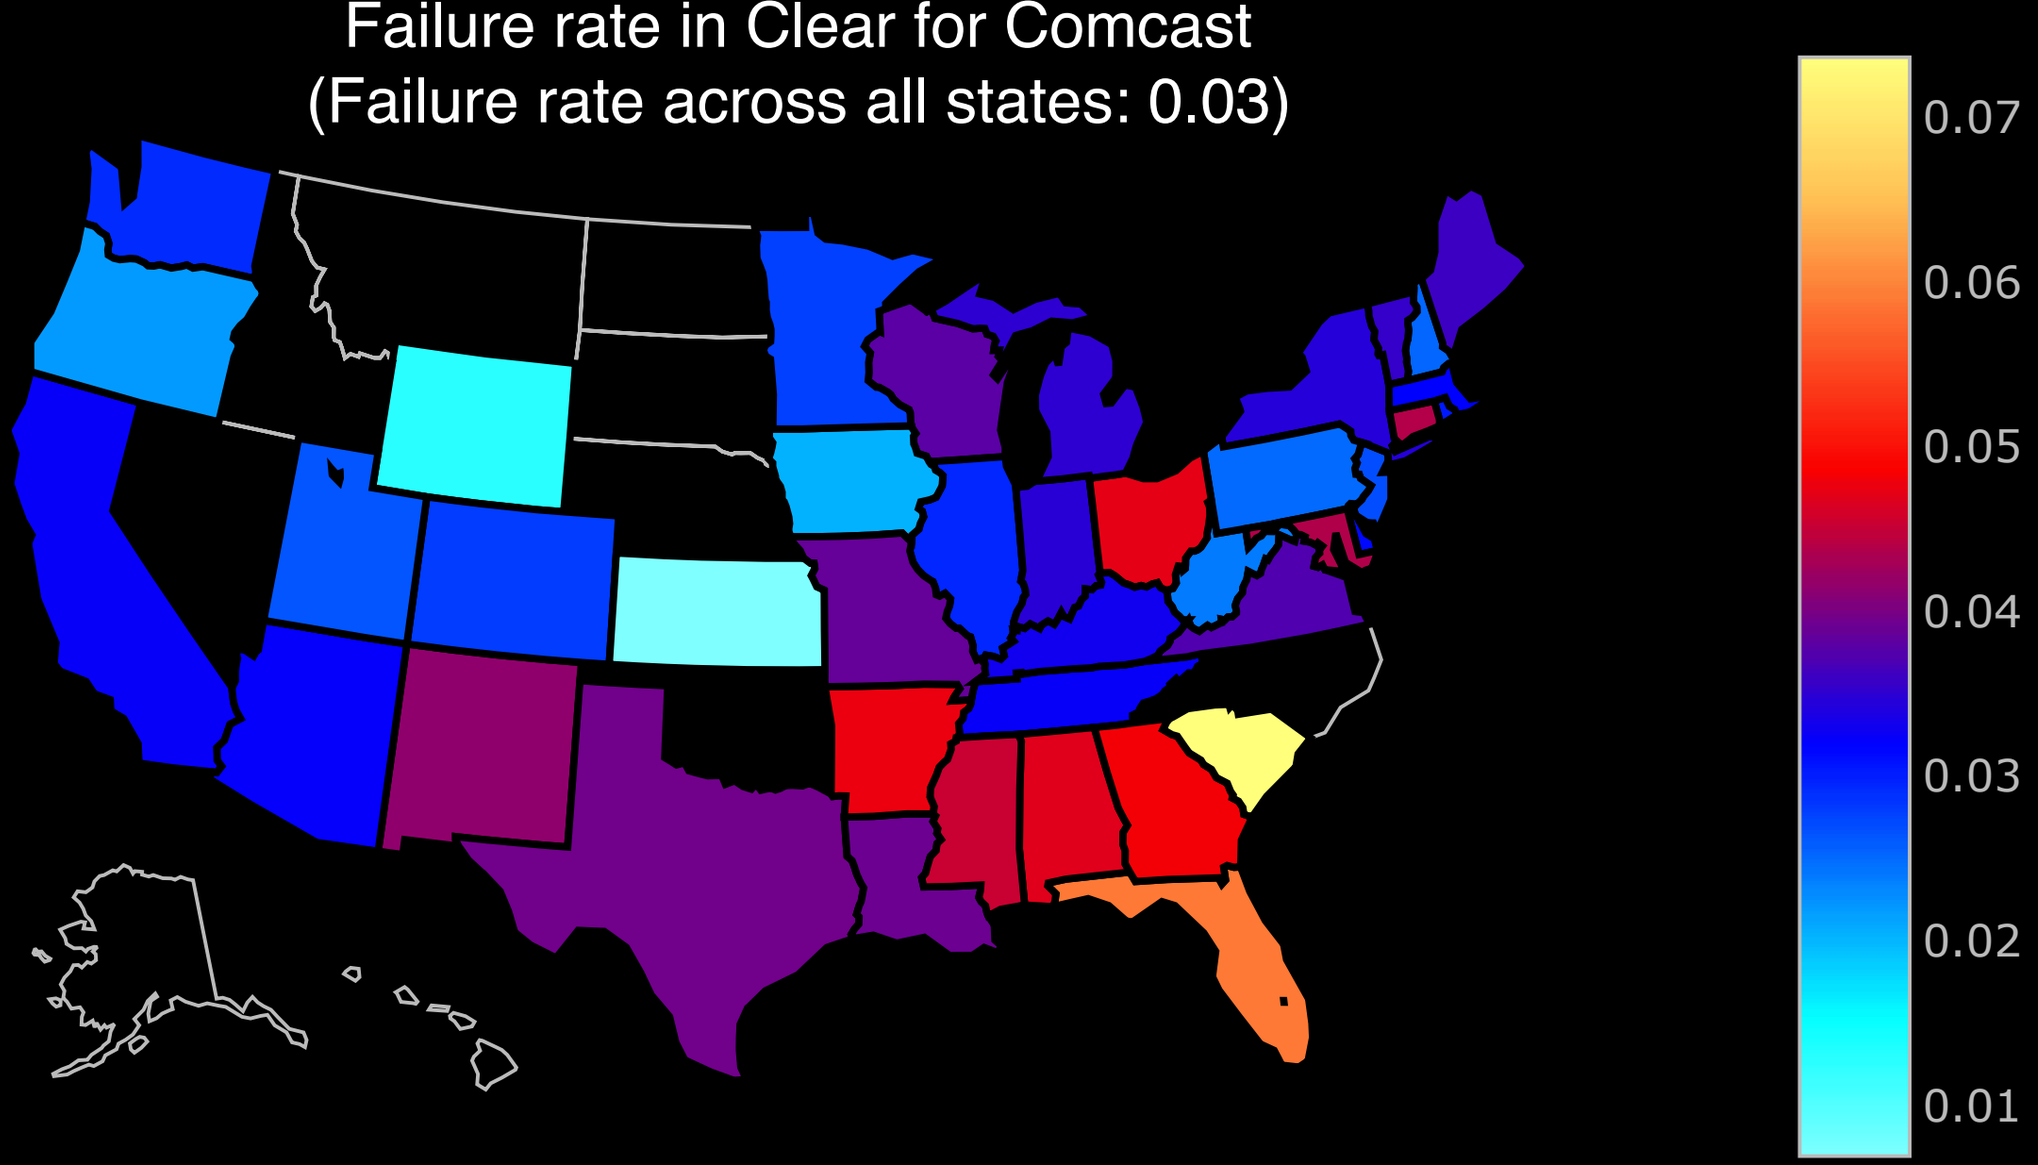

For Centurylink, failure rate in clear weather is higher in the west

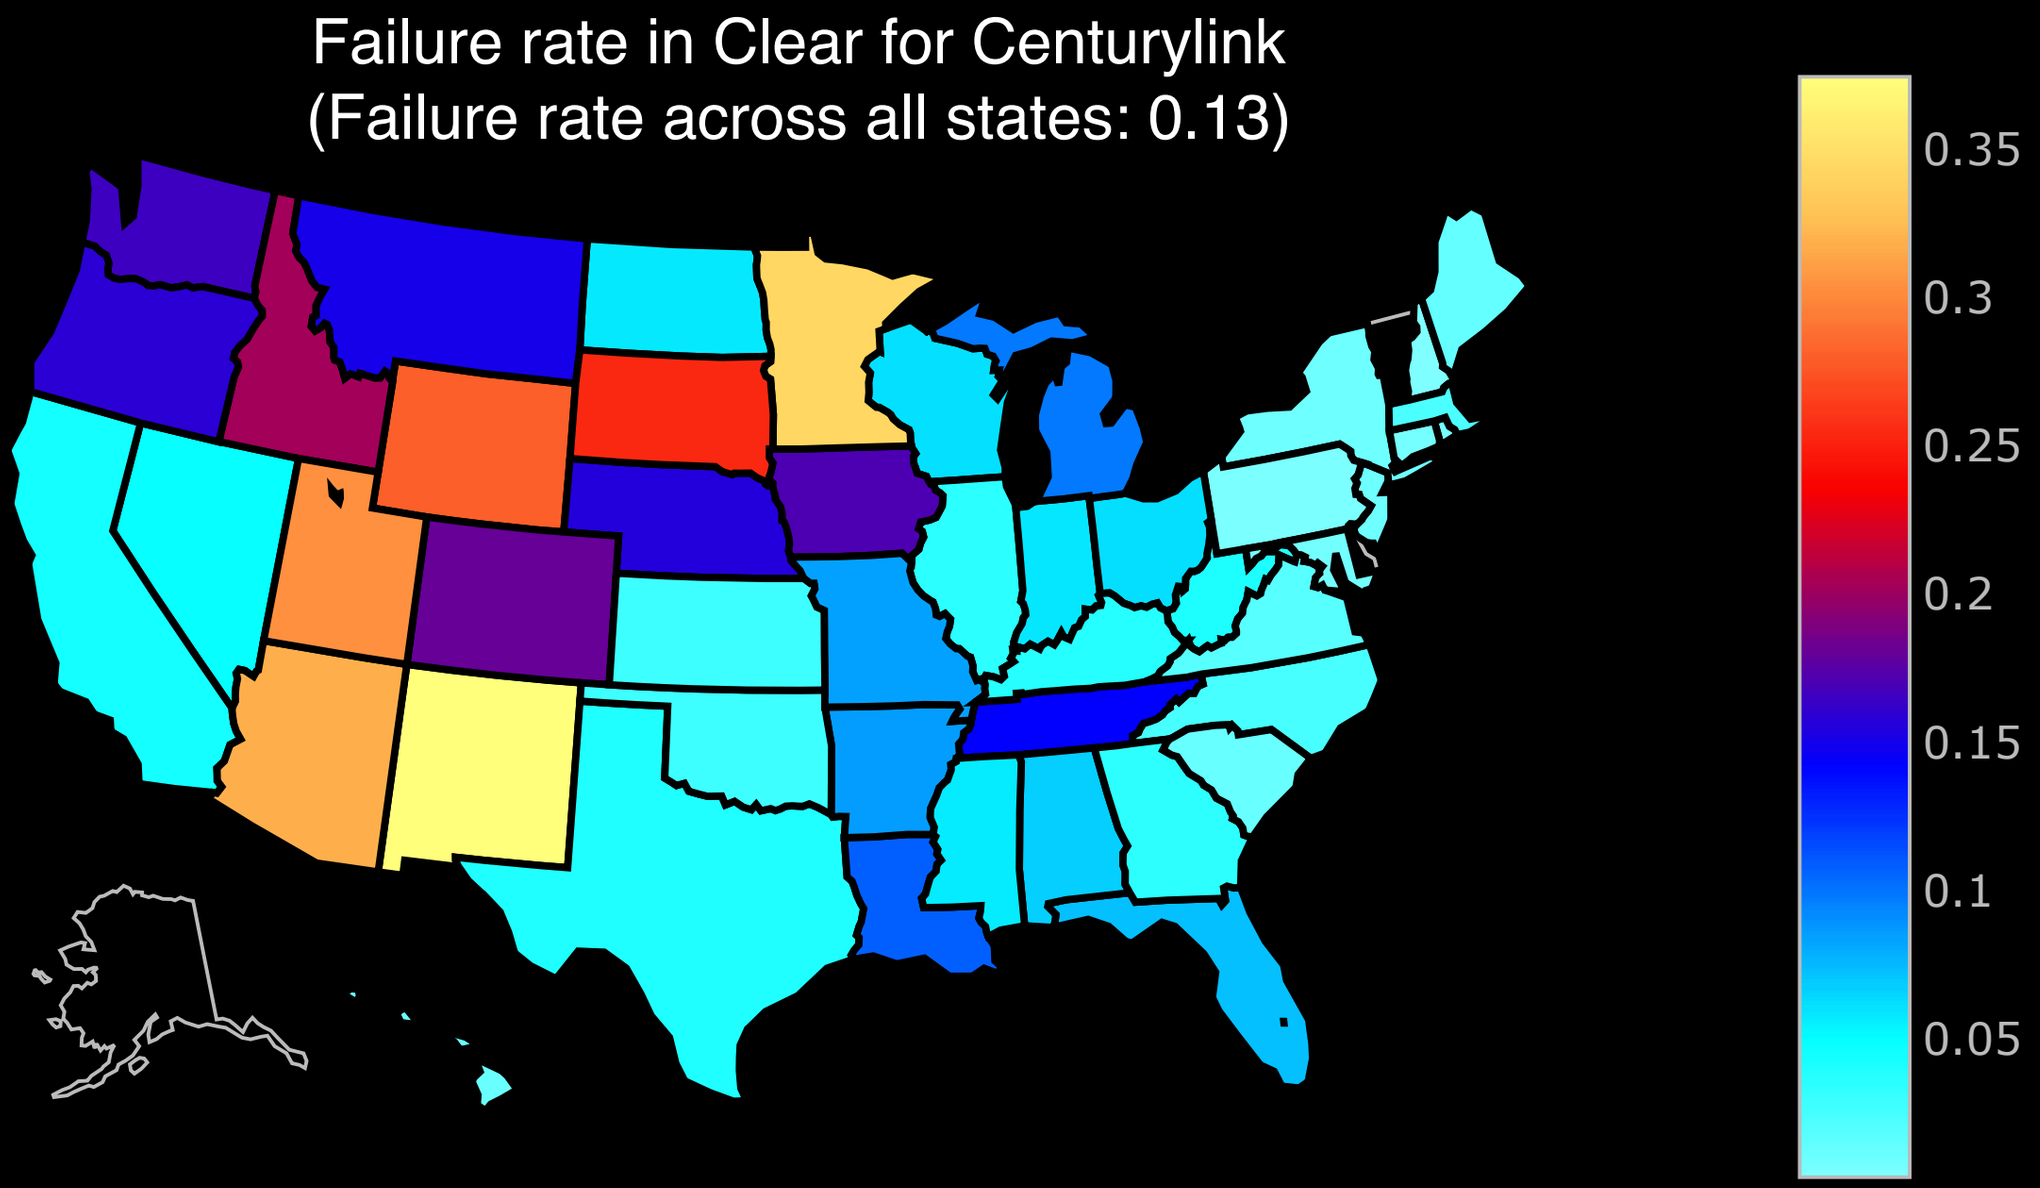

For Charter, failure rate in clear weather is more uniform

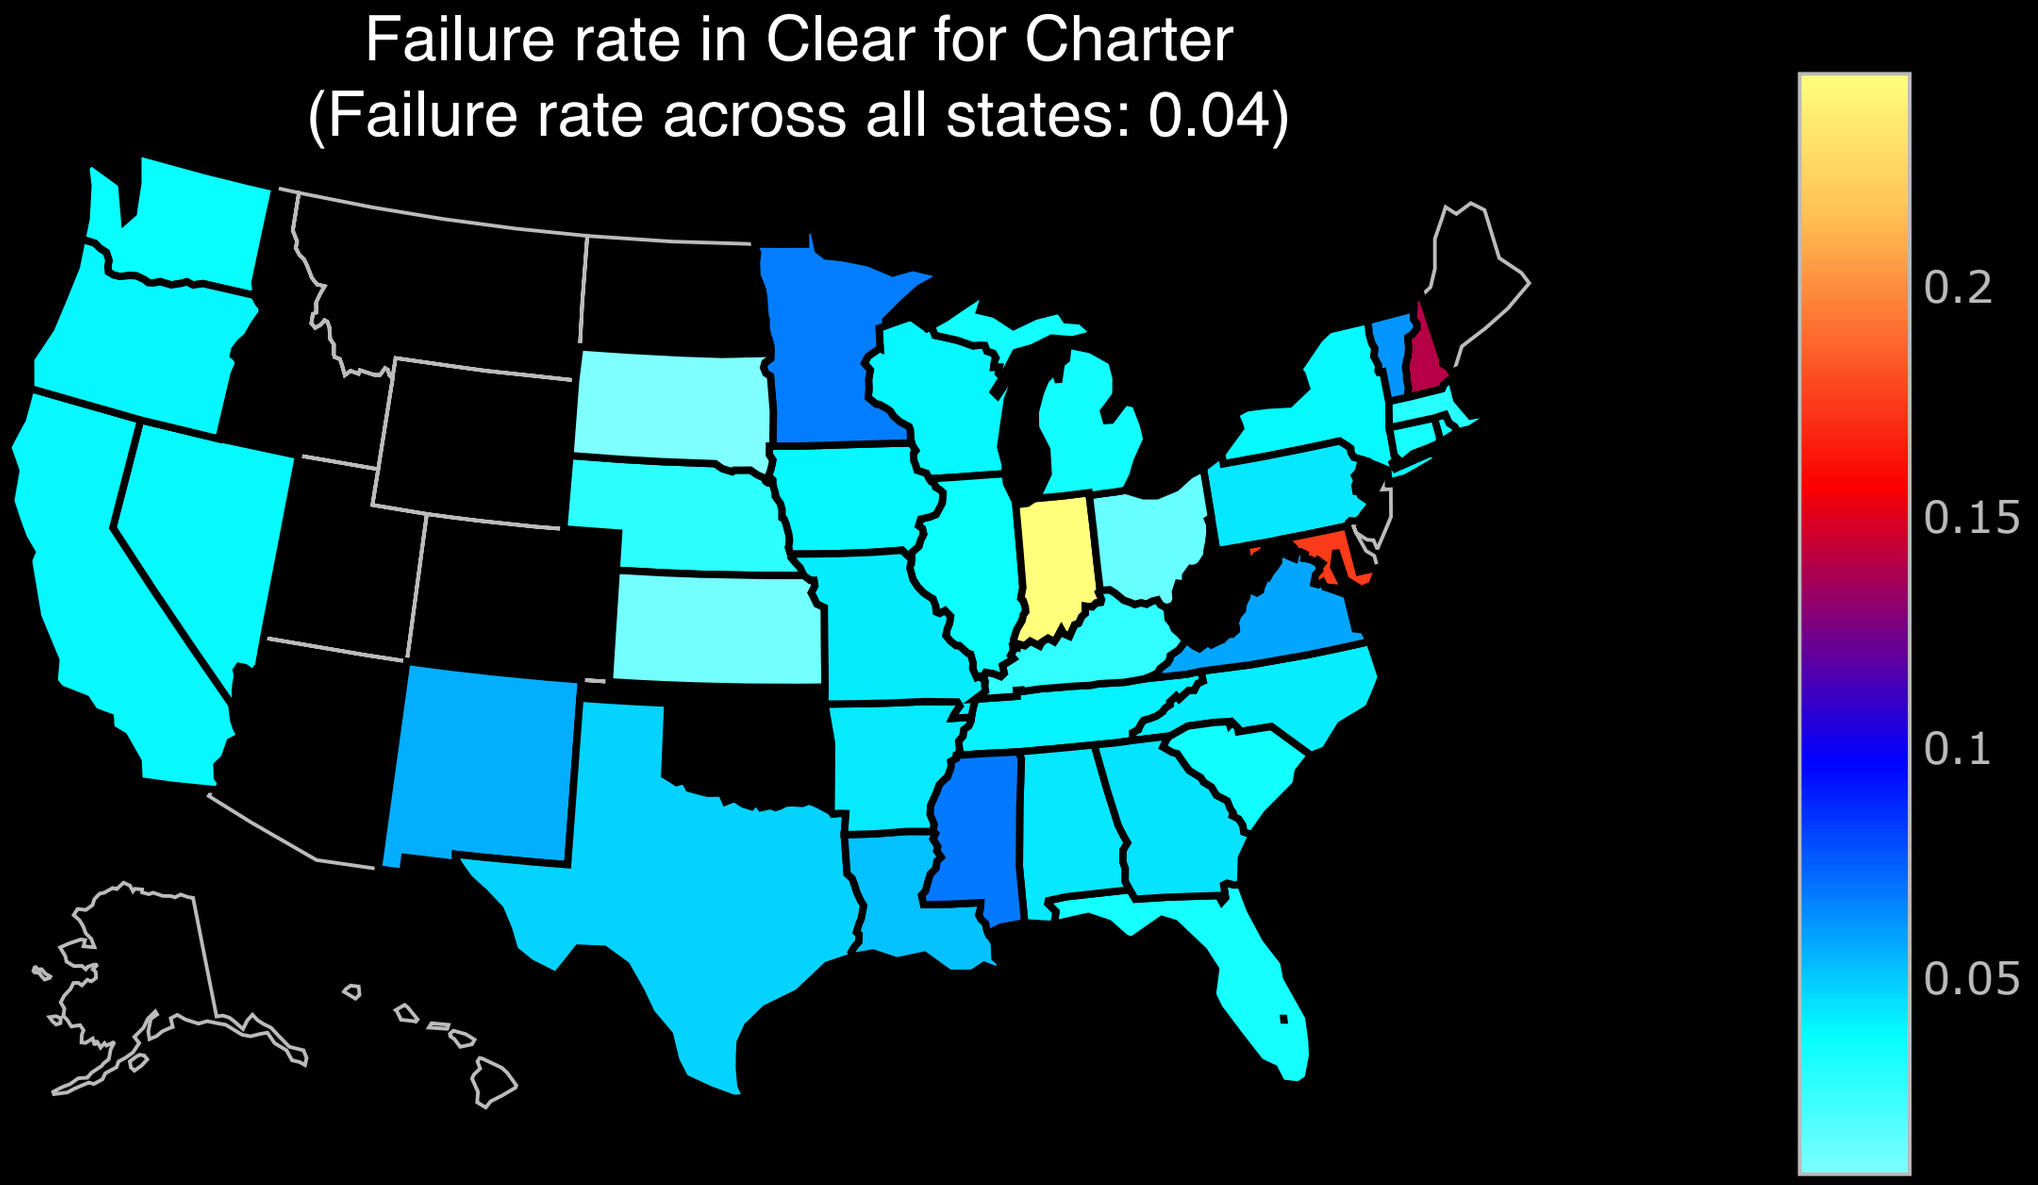

Residential Internet infrastructure reliability varies across regions

- Failure rates vary across U.S. states even in clear weather
  - Snow had higher failure rates in the southern states (see backup slides)
- Upcoming work
  - Longitudinal analysis: Are trends observed in 2017 also present in earlier years?
  - Large events: Can we use correlated failures to detect events such as power outages?

Backup Slides

Study reliability across weather conditions using normalized failure rates

- Find failure rate in Clear
- Find failure rate in weather condition
- Calculate normalized failure rate as failure-rate-in-weather/failure-rate-in-clear

Failure rates in snow, rain, and thunderstorm, are higher than in clear

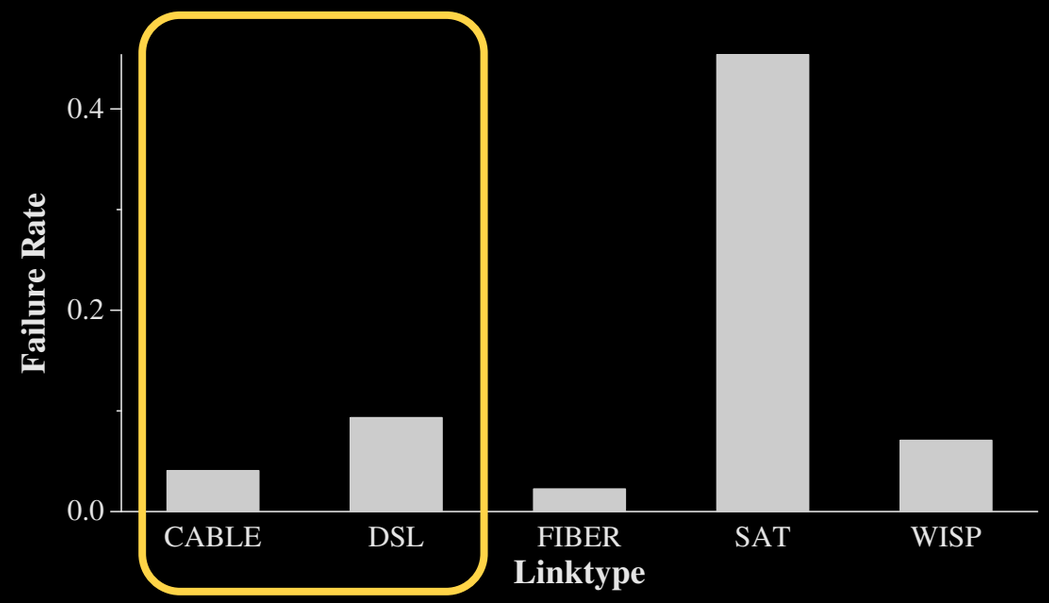

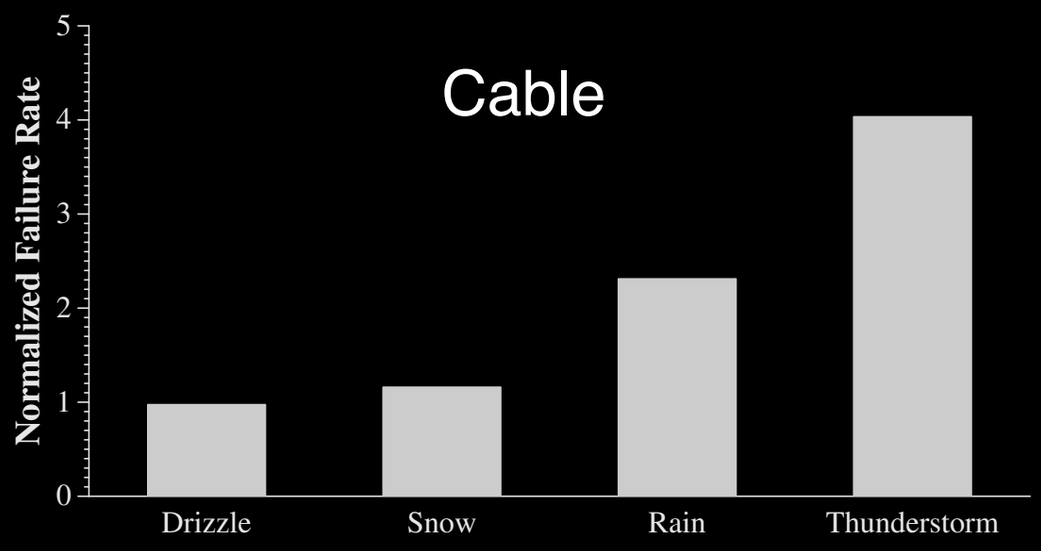

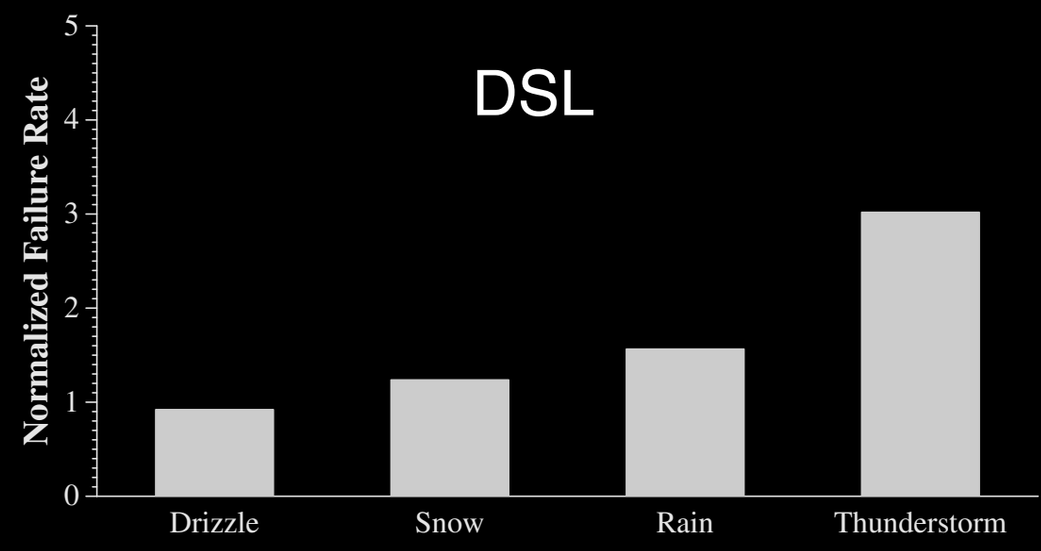

For Cable, the south experiences higher failure rates in Snow

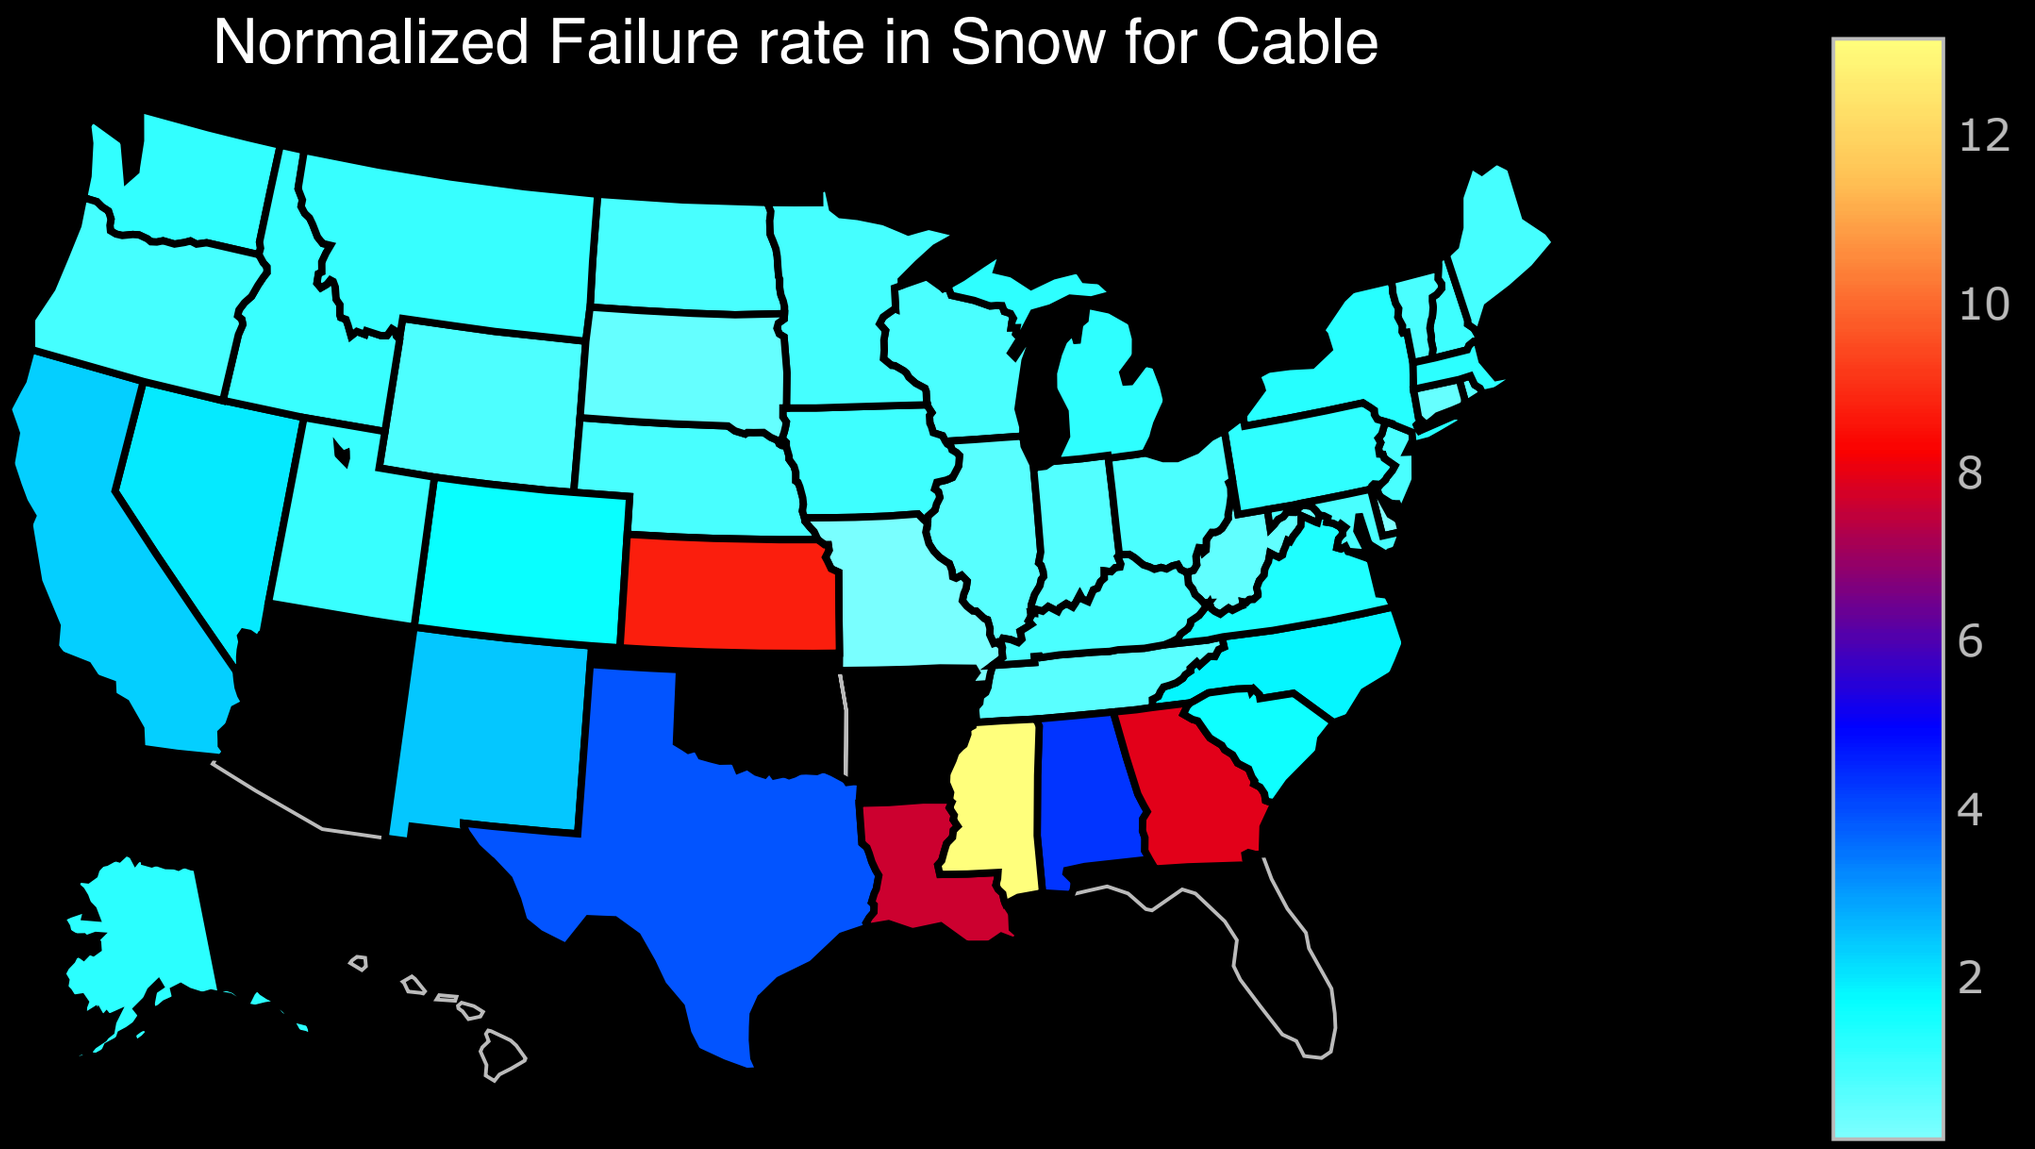

DSL also has higher failure rates in snow in the south

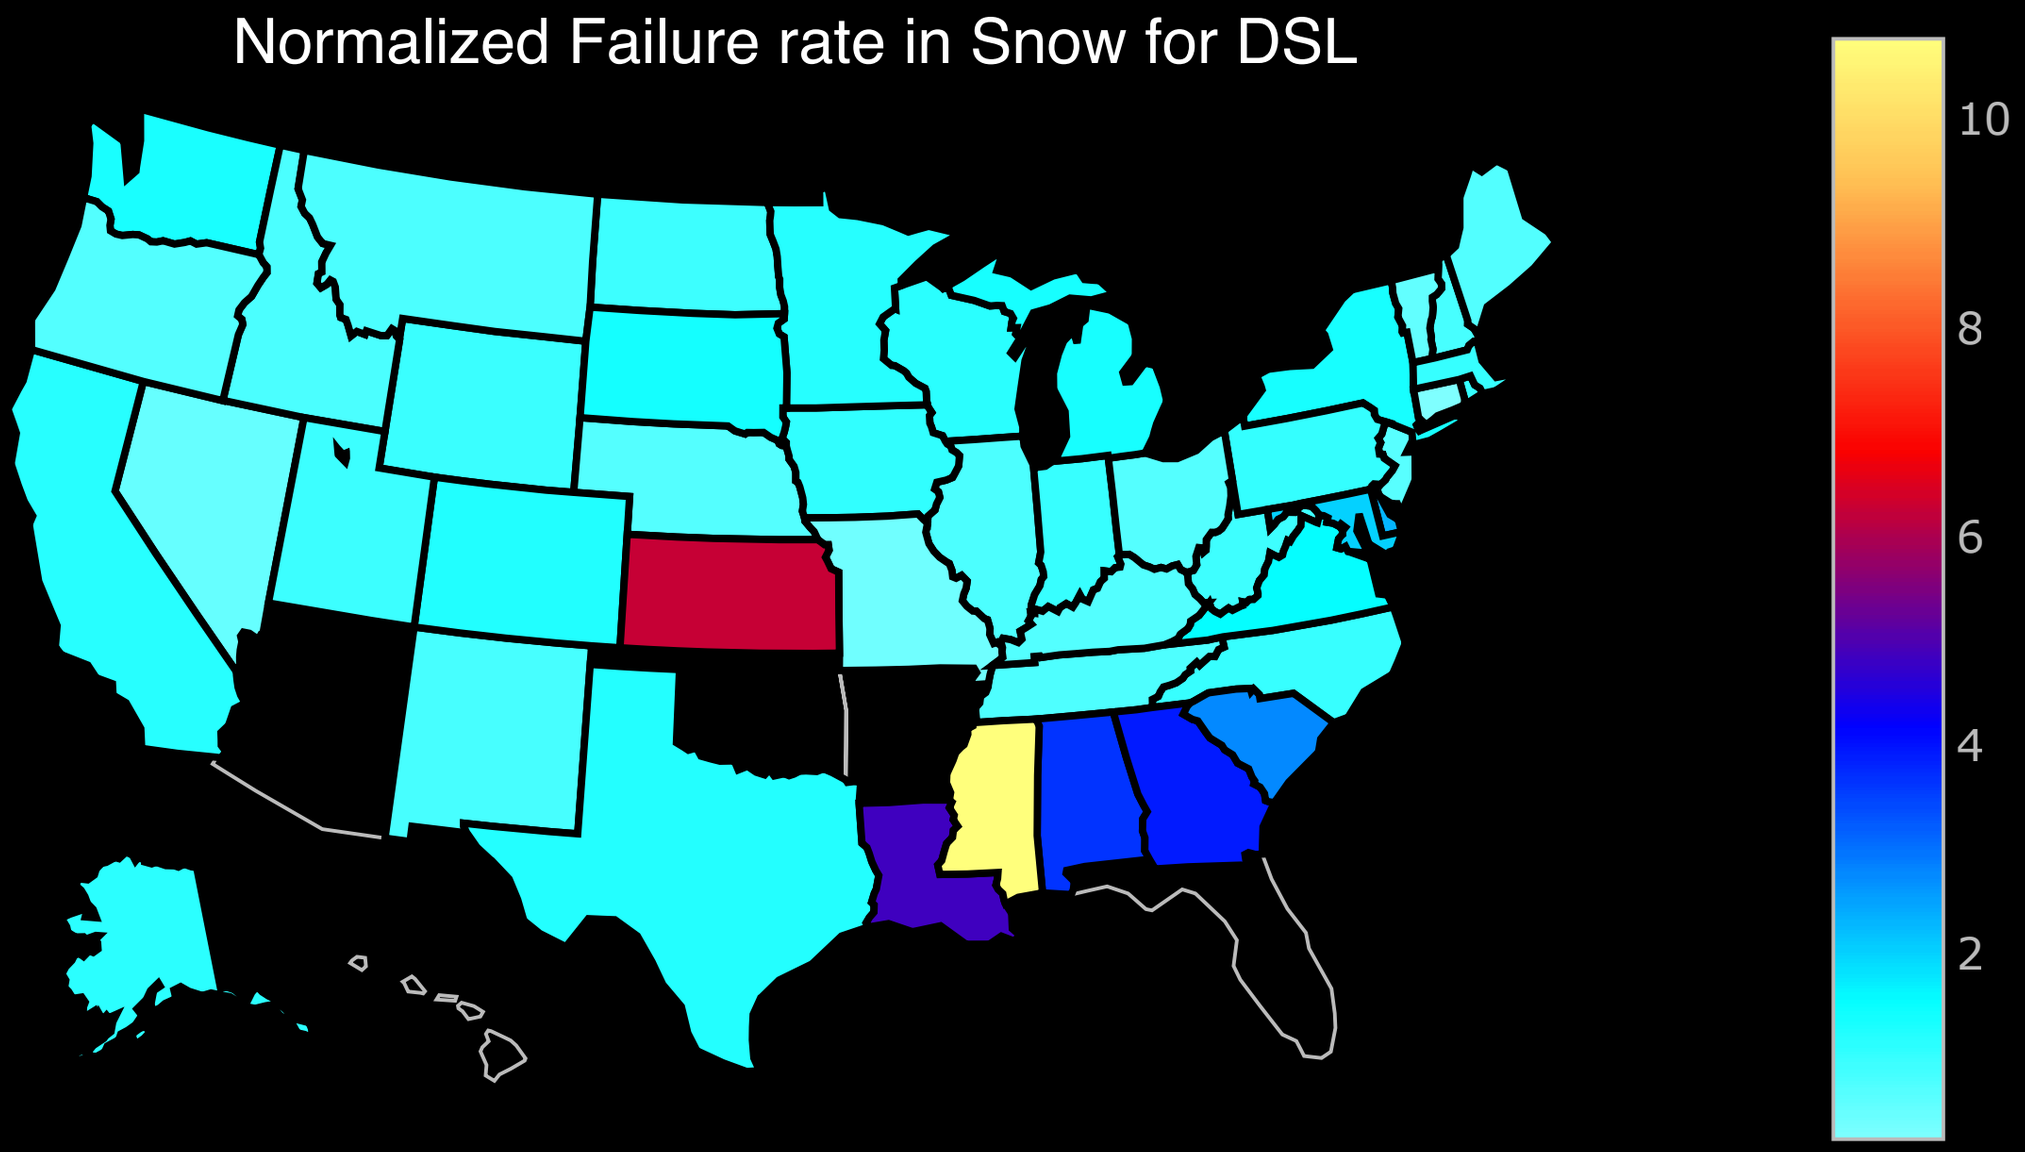

Cable has higher failure rates in the north east and south east in rain

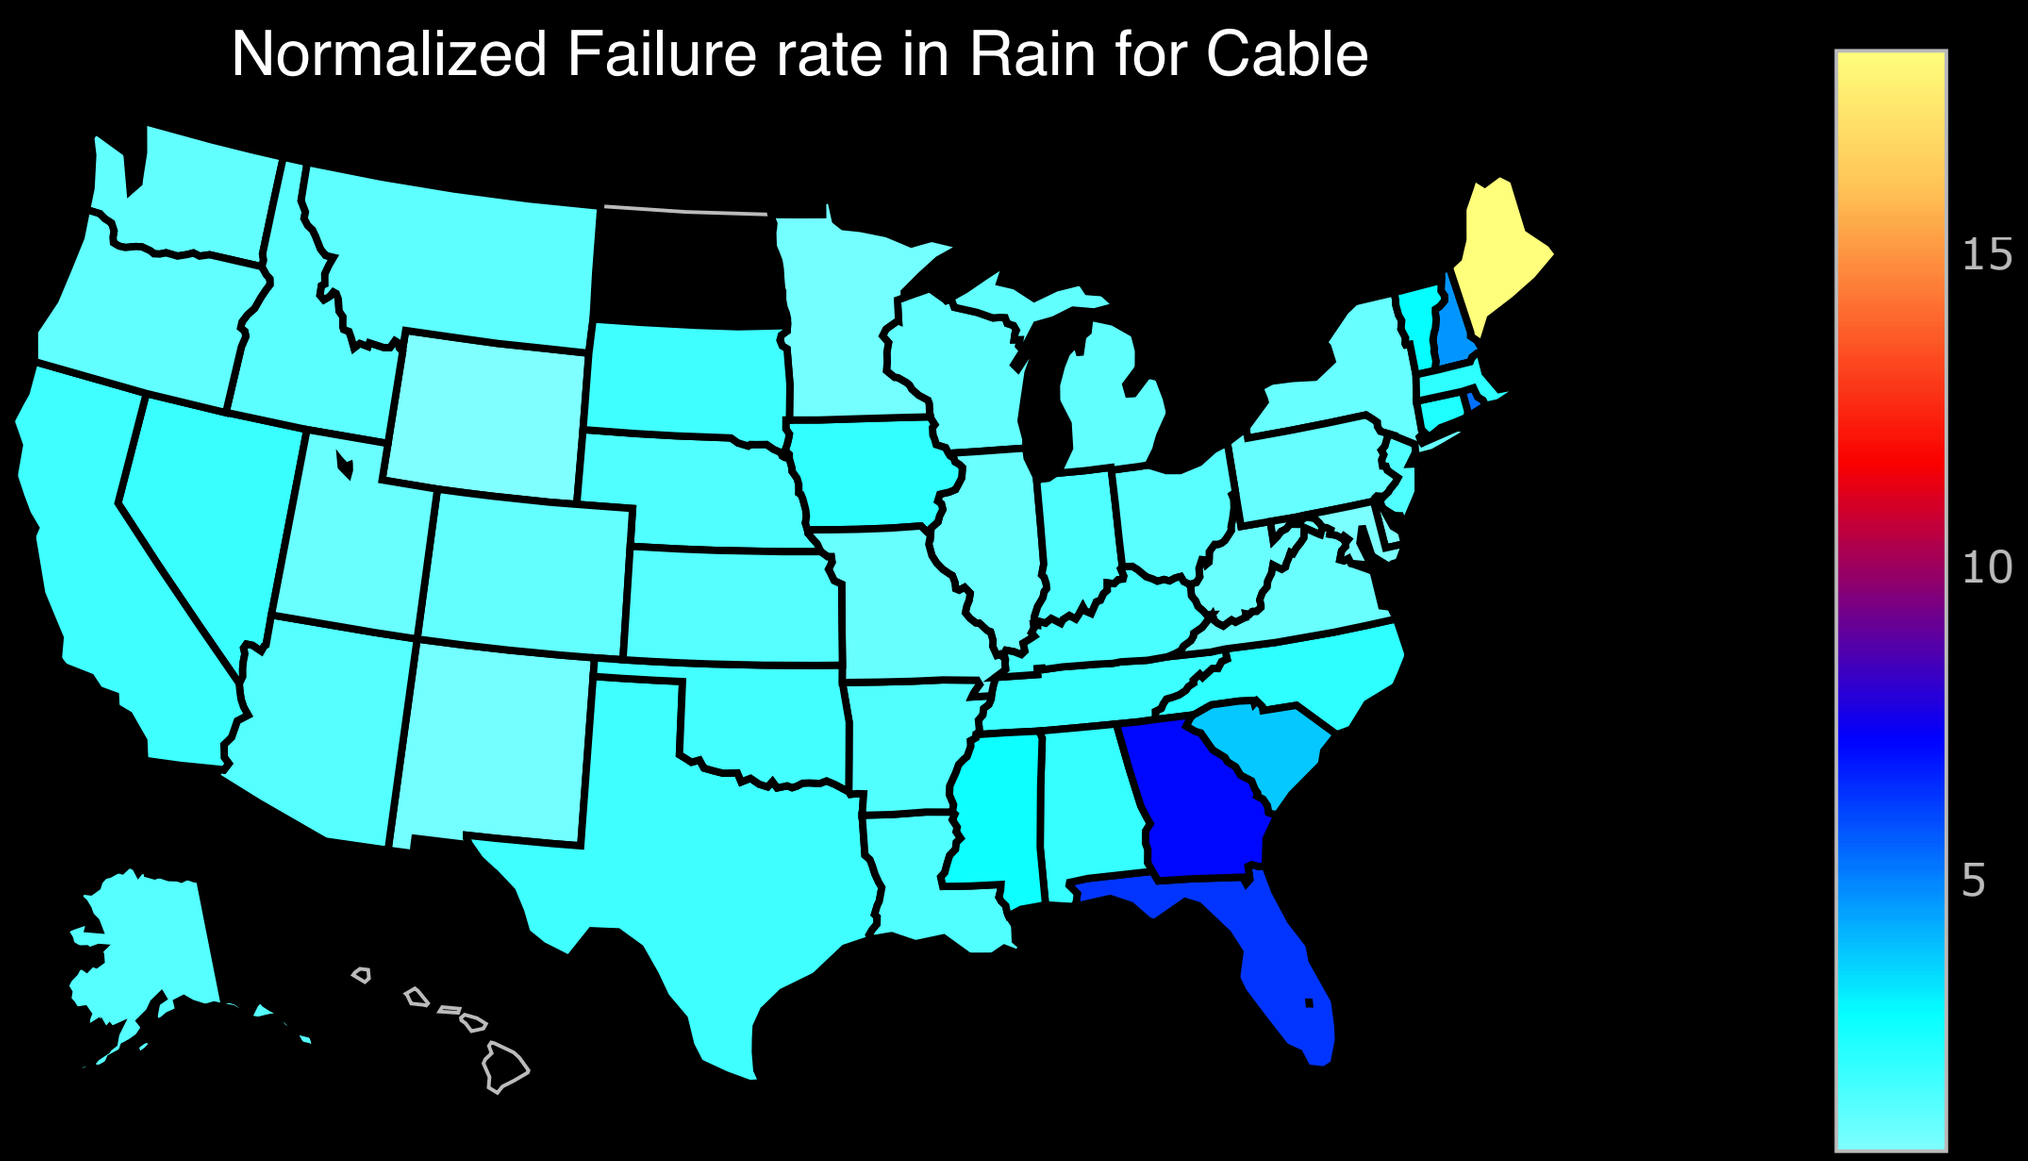

As does DSL, though the increase in failure rate is less pronounced