

# A Comparison of Syslog and IS-IS for Network Failure Analysis



**Daniel Turner**  
Kirill Levchenko  
Stefan Savage  
Alex C. Snoeren


# Network Reliability

- Networks component failure at scale is inevitable



# Network Reliability

- ◆ Networks component failure at scale is inevitable
- ◆ Many mechanisms in place to keep customers from noticing
  - ◆ Redundant hardware & protocols



# Network Reliability

- ◆ Networks component failure at scale is inevitable
- ◆ Many mechanisms in place to keep customers from noticing
  - ◆ Redundant hardware & protocols
- ◆ Evaluating reliability mechanisms requires data


# Network Reliability Data

- ◆ Syslog has been popular for this role
  - ◆ Easy to obtain and utilize (open source & commercial tools)
  - ◆ [Gill Sigcomm11], [Mahimkar Sigcomm09], [Qiu IMC10], [Potharaju Sigmetrecs13], [**Turner Sigcomm10**]
- ◆ The **gold** standard is direct IGP routing messages capture
  - ◆ Fate sharing with network
  - ◆ Less widely used because its harder to obtain

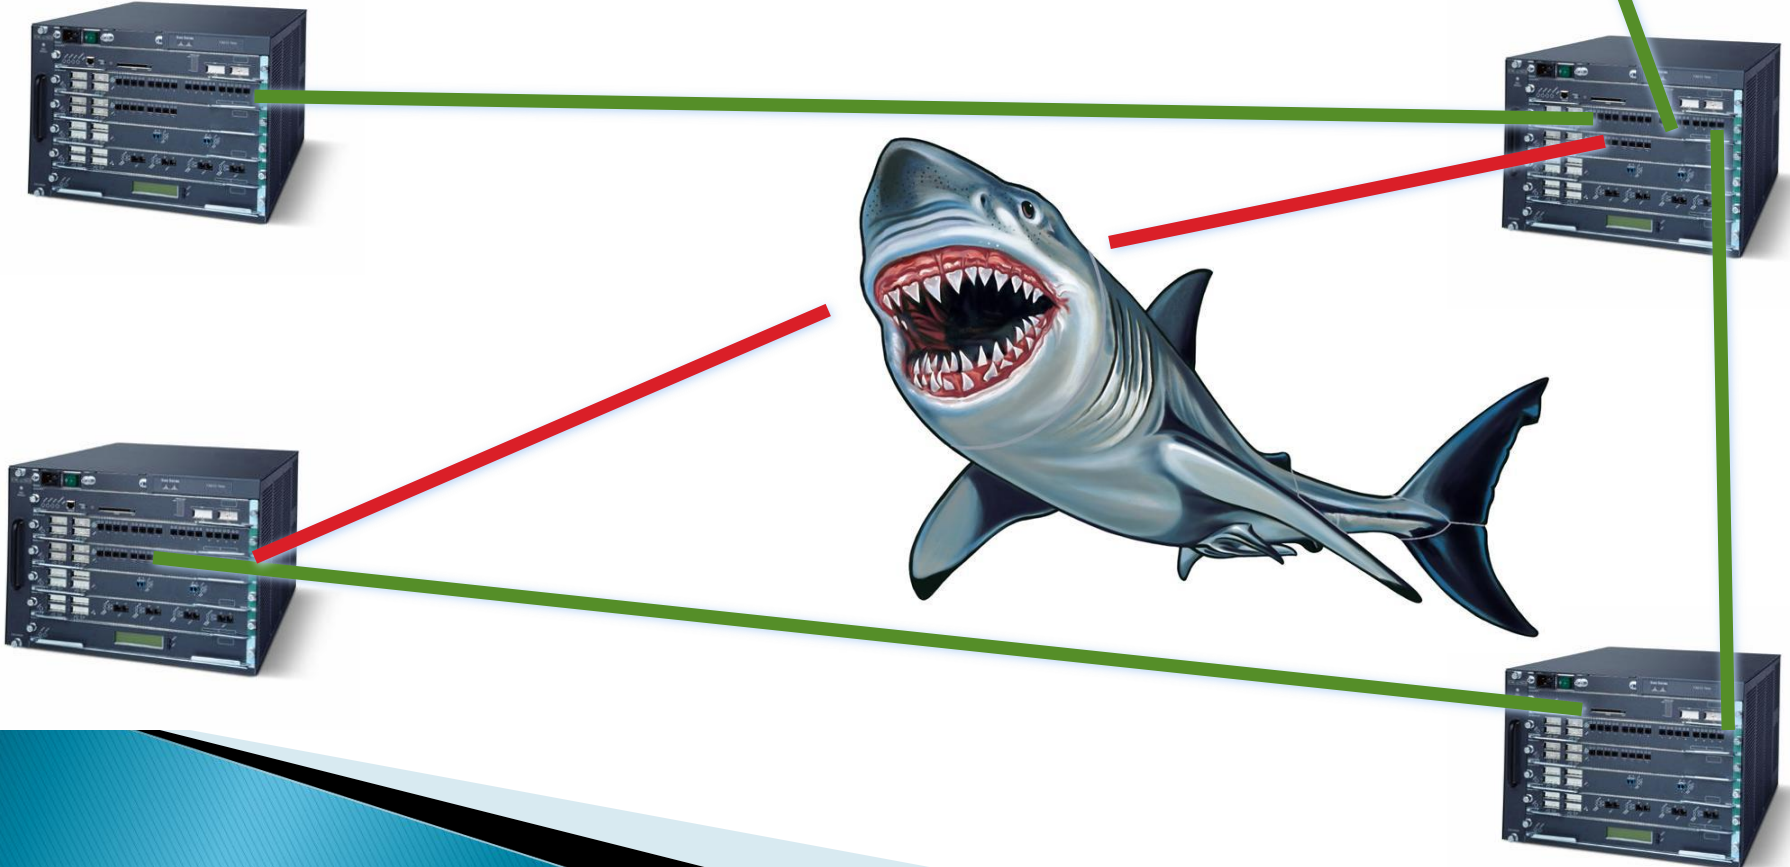

# Failure Example





# Failure Example



# Failure Example



# Failure Example





Router x:  
Interface 1/1  
DOWN



# Failure Example



IS-IS  
LSP



# Failure Example

## IS-IS Link State Packet

ID: Router 3

Time: 2/2/11 3:00PM

Current Neighbors:

\* Router 4 : weight 27

...



# Data Usage

- ◆ **How accurate is syslog, as compared to IS-IS, when used to capture and characterize failure?**
- ◆ Different actors have different needs from the data
  - ◆ Details about root cause
  - ◆ Frequency and duration
  - ◆ Failure impact

# Comparative Analysis

- ◆ Question 1: Can syslog be used as a drop in replacement for IS-IS data?
- ◆ Question 2: For what purposes can syslog be used as a replacement for IS-IS data?
- ◆ Question 3: If you are limited to only using syslog what can be done to improve its accuracy?

# Data Collection

## ◆ CENIC Network

- ◆ ISP to California educational institutions
- ◆ 225+ routers
- ◆ 299 Links
- ◆ Thousands of miles of fiber
- ◆ Millions of daily users

## ◆ 13 Months of data

- ◆ 11 Million IS-IS LSPs
- ◆ 47,000 Syslog Messages

# Syslog as a drop in replacement

- ◆ What is required to be a drop in replacement?
  - ◆ State of the network as seen by both data sources is the same
- ◆ We are focusing on link state (Up / Down)
  - ◆ Function of state transitions
- ◆ Do syslog's state transitions mirror IS-IS's?
  - ◆ Straightforward to measure and compare

# Examining State Transitions

Transitions	Router Syslog Messages		
	None	One	Both
DOWN			4,962 (43%)
UP			

# Examining State Transitions

Transitions	Router Syslog Messages		
	None	One	Both
DOWN		4,512 (39%)	4,962 (43%)
UP			


# Examining State Transitions

Transitions	Router Syslog Messages		
	None	One	Both
DOWN	2,022 (18%)	4,512 (39%)	4,962 (43%)
UP			

# Examining State Transitions

Transitions	Router Syslog Messages		
	None	One	Two
DOWN	2,022 (18%)	4,512 (39%)	4,962 (43%)
UP			

18% is huge



# Examining State Transitions

Transitions	Router Syslog Messages		
	None	One	Two
DOWN	2,022 (18%)	4,512 (39%)	4,962 (43%)
UP	1,696 (15%)	5,432 (48%)	4,168 (37%)

18% is huge

# What are the implications

- ◆ Syslog is not a drop in replacement for IS-IS data
  - ◆ Can't do failure for failure accounting
- ◆ Question 2: For what purposes can syslog be used as a replacement for IS-IS data?
- ◆ Some people only need statistical similarity
  - ◆ Statistics are usually about failures not state changes

# Link Failures

	IS-IS	Syslog	Overlap
Downtime (Hours)	3,648	2,714	2,331
Failure Count	11,213	11,738	9,298

# Link Failure

Missing 1k hours of  
downtime

	IS-IS	Syslog	Overlap
Downtime (Hours)	3,648	2,714	2,331
Failure Count	11,213	11,738	9,298

# Link Failure

20% of syslog failures are false positives

Missing 1k hours of downtime

	IS-IS	Syslog	Overlap
Downtime (Hours)	3,648	2,714	2,331
Failure Count	11,213	11,738	9,298

# What are the implications

- ◆ Not all statistics will match
  - ◆ But some could
- ◆ Statistical similarity measured w/ Komogorov-Smirnov test

	Backbone Links	Customer Access Links
Annualized Downtime per Link		
Annualized Failures per Link		
Failure Duration		

# What are the implications

- ◆ Not all statistics will match
  - ◆ But some could
- ◆ Statistical similarity measured w/ Komogorov-Smirnov test

	Backbone Links	Customer Access Links
Annualized Downtime per Link	Yes	Yes
Annualized Failures per Link	Yes	Yes
Failure Duration	No	No

# What are the implications

- Not all statistics will match
  - But some could
- Statistical similarity measured w/ Komogorov-


Due to false positives and can be fixed

	Backbone Links	Customer Access Links
Annualized Downtime per Link	Yes	Yes
Annualized Failures per Link	Yes	Yes
Failure Duration	No	No

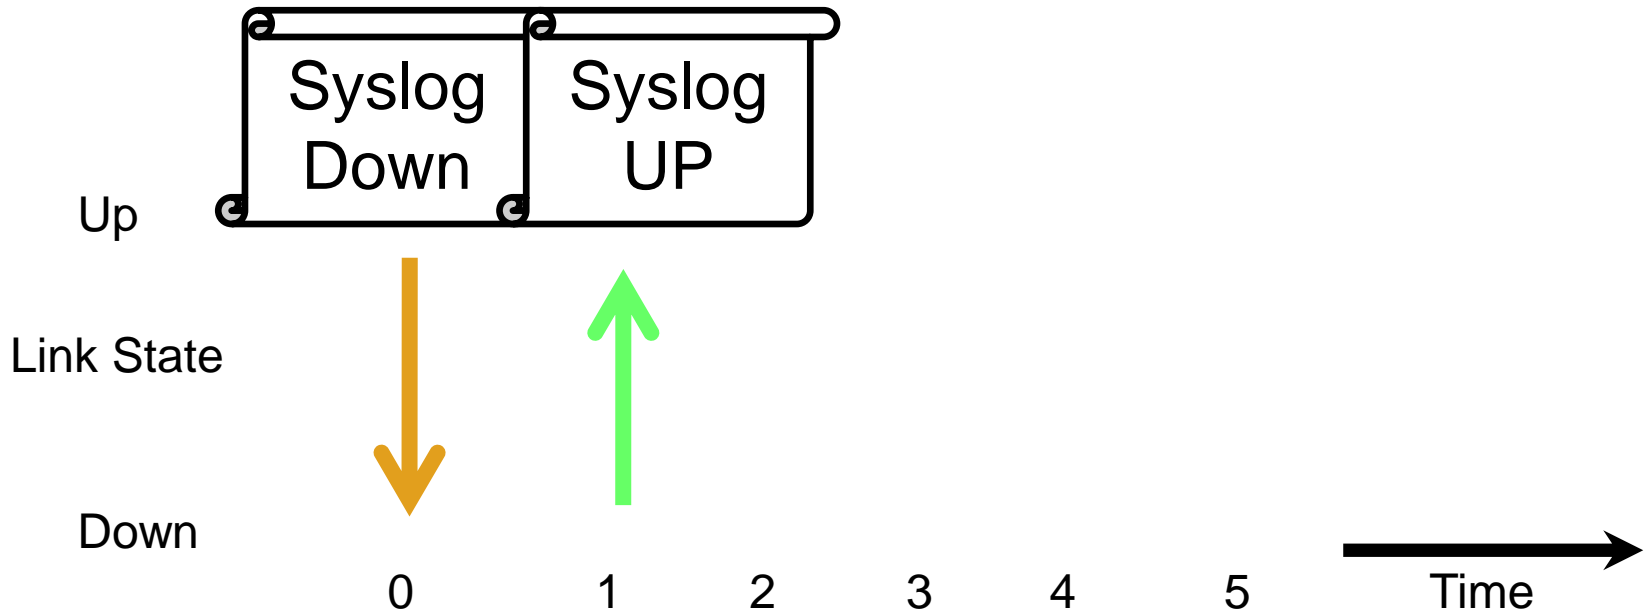
# Improving Syslog's Fidelity

- ◆ Question 3: If you are limited to only using syslog what can be done to improve its accuracy?
  - ◆ Eliminate false positives
    - ◆ Mostly very short failures
  - ◆ Remove ambiguous state transitions
    - ◆ 8% of link time is between to ambiguous transitions
- ◆ How do we know this?
  - ◆ We have access to both syslog and IS-IS data


# Ambiguous State Transitions




# Ambiguous State Transitions




# Ambiguous State Transitions




# Ambiguous State Transitions



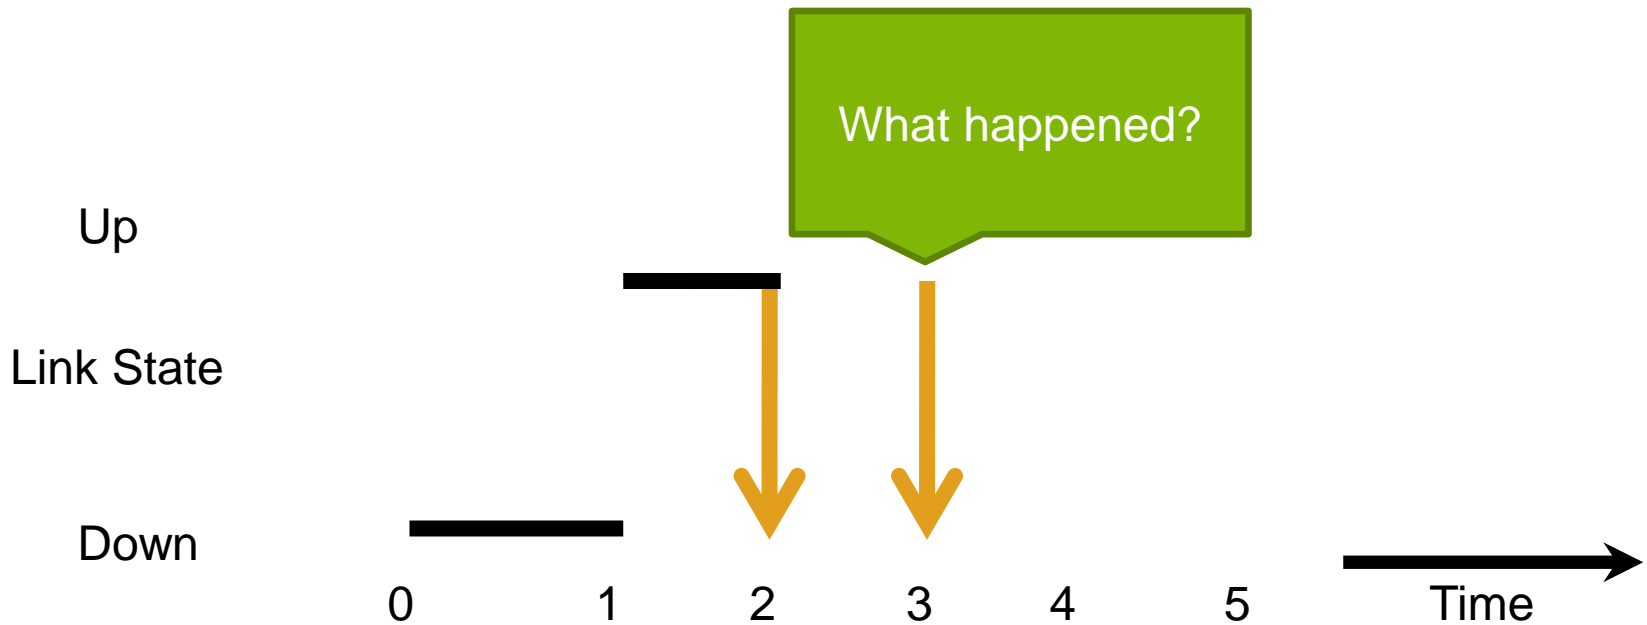
# Ambiguous State Transitions



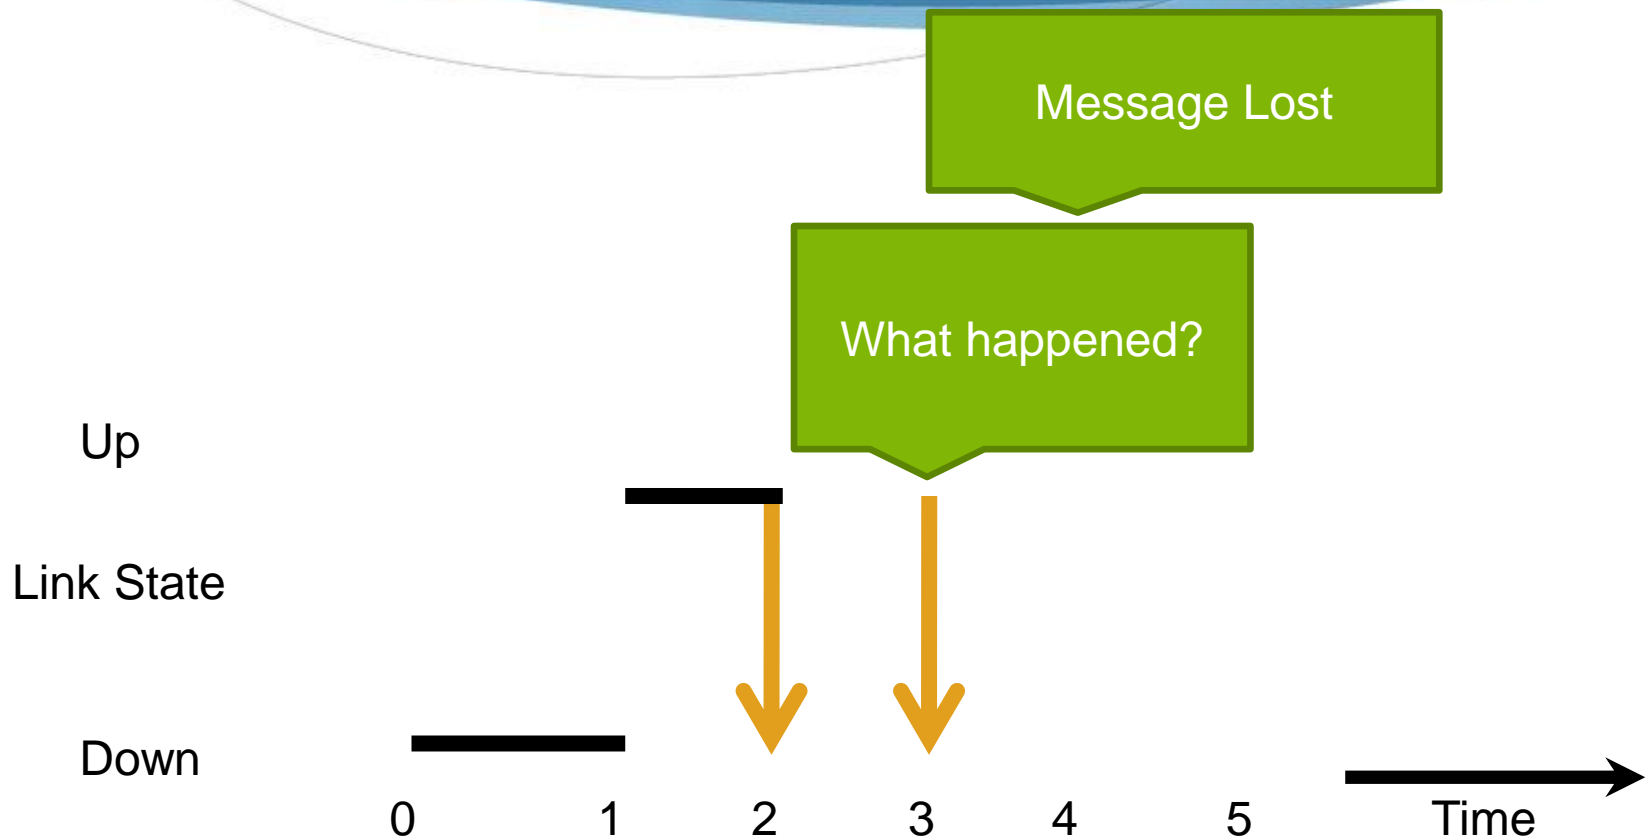
# Ambiguous State Transitions



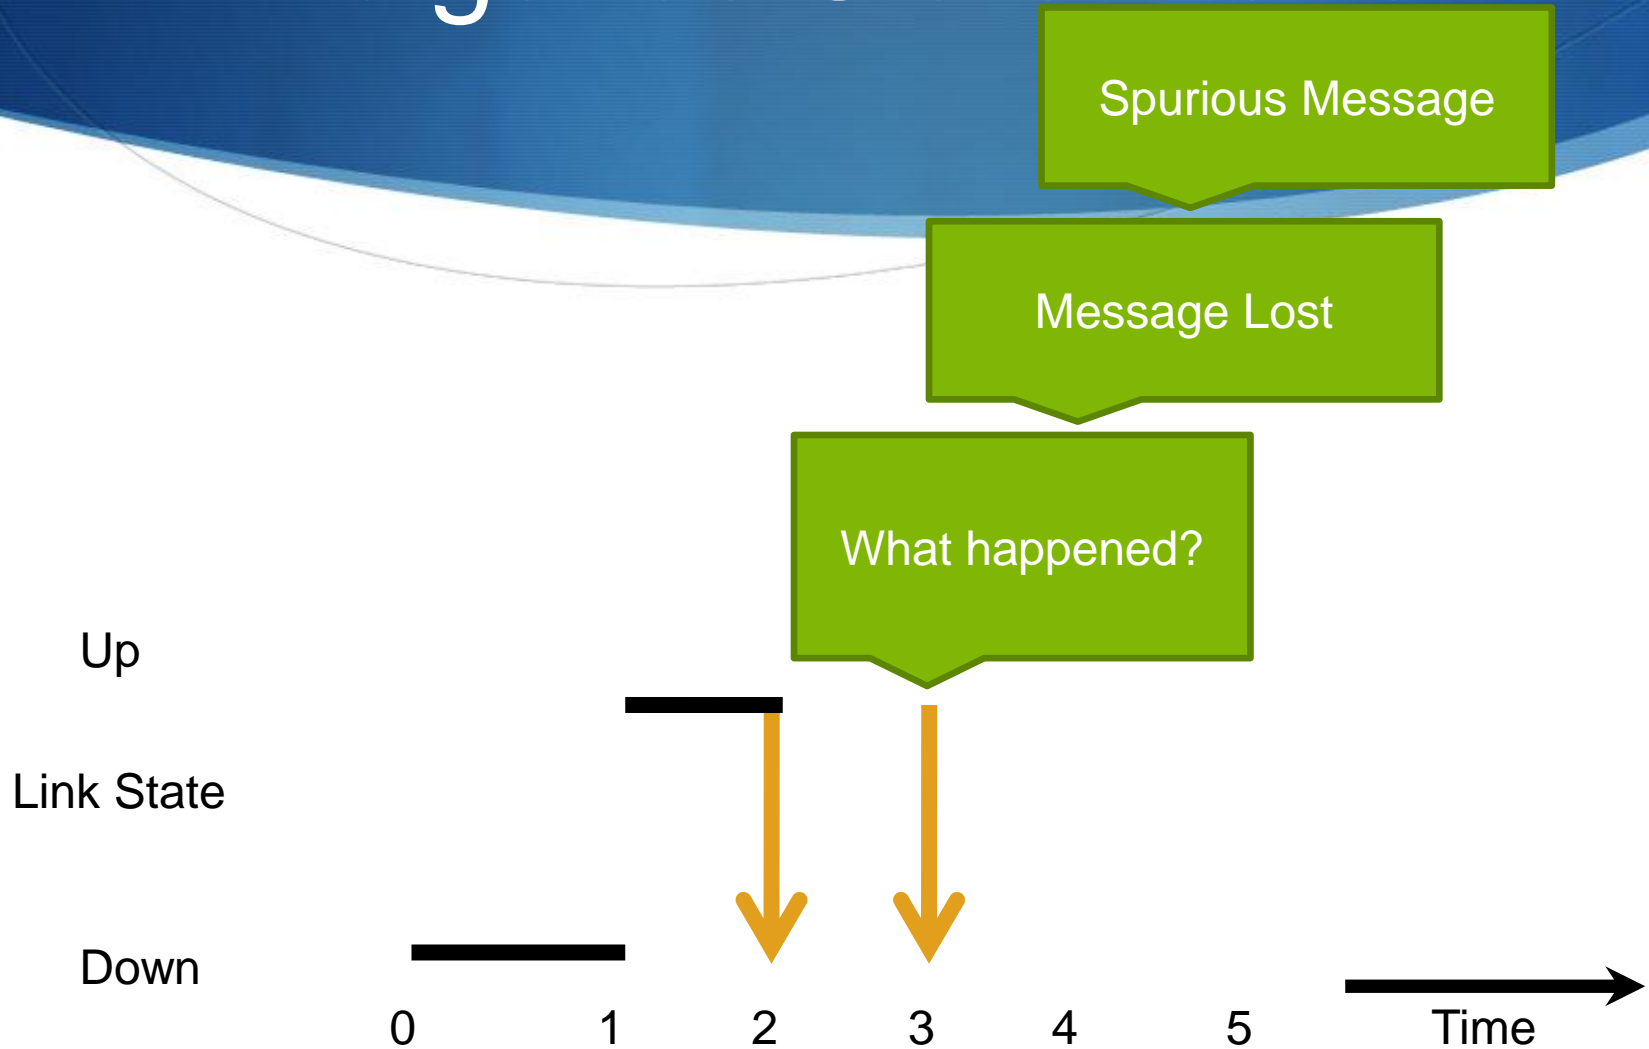
# Ambiguous State Transitions



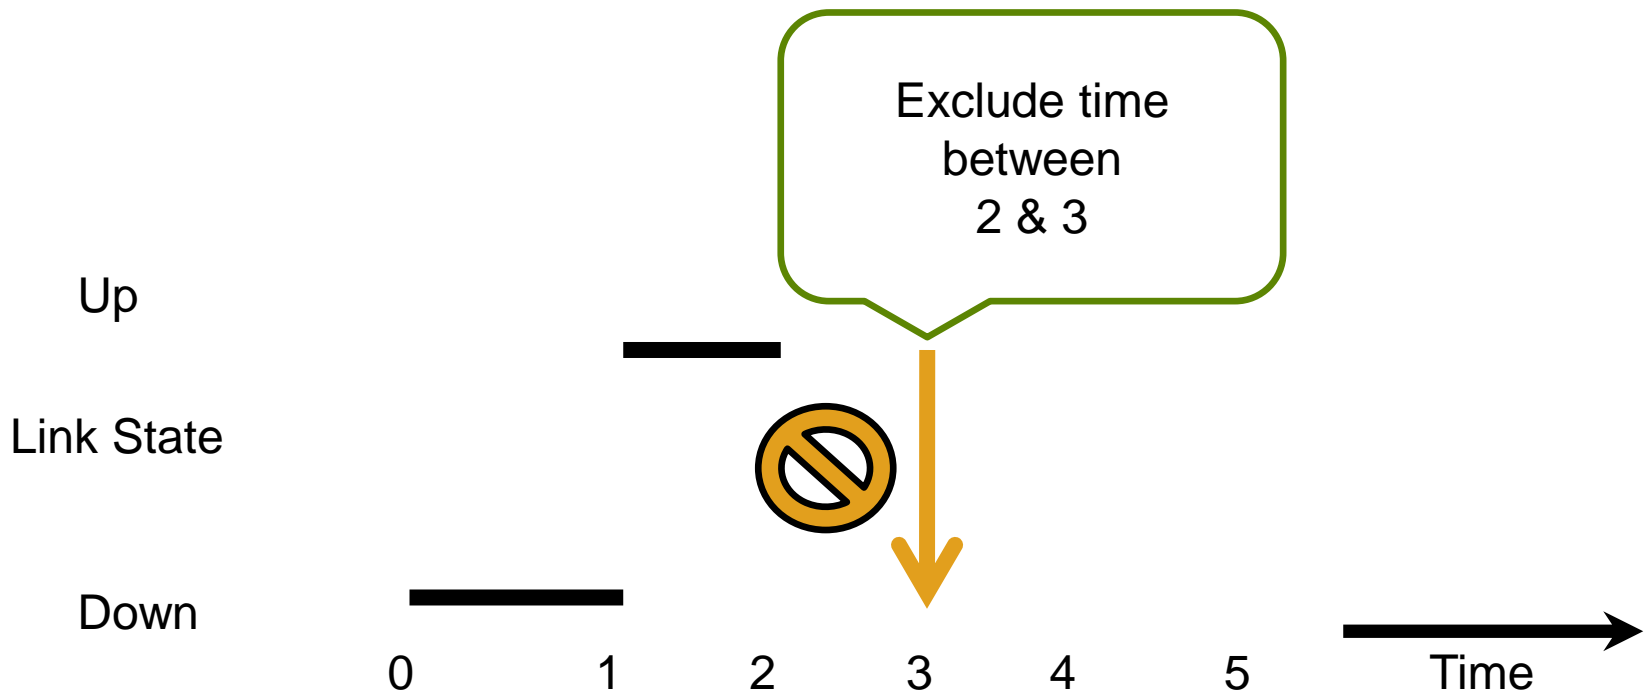
# Ambiguous State Transitions




# Ambiguous State Transitions




# Ambiguous State Transitions




# Ambiguous State Transitions



# Ambiguous State Transitions



# Ambiguous State Transitions



# Correcting Ambiguous Transitions

- ◆ Strategies to best improve syslog's fidelity
  - ◆ Always down? Always up?
  - ◆ Ignore the first? Ignore the second?

# Correcting Ambiguous Transitions

- ◆ Strategies to best improve syslog's fidelity
  - ◆ Always down? Always up?
  - ◆ Ignore the first? Ignore the second?

	DOWN	UP
Lost Message	42%	86%
Spurious retransmit	52%	14%
Other	6%	0%

# Correcting Ambiguous Transitions

- 💧 Strategies to best improve syslog's fidelity
  - 💧 Always down? Always up?
  - 💧 Ignore the first? Ignore the second?

Almost always multiple retransmits per failure

	DOWN	UP
Lost Message	42%	86%
Spurious retransmit	52%	14%
Other	6%	0%

# Correcting Ambiguous Transitions

- ◆ Strategies to best improve syslog's fidelity
  - ◆ Always down? Always up?
  - ◆ Ignore the first? Ignore the second?
- ◆ Optimal strategy: ignore the second message

	DOWN	UP
Lost Message	42%	86%
Spurious retransmit	52%	14%
Other	6%	0%

# Conclusion

- ◆ Syslog not a drop in replacement for IGP data when studying failure
- ◆ Can be used to measure aggregate failure characteristics
  - ◆ Downtime & Failure counts
- ◆ Filtering can improve syslog's fidelity
- ◆ When possible we recommend setting up an IGP listener

