On The Fidelity of 802.11 Packet Traces

Aaron Schulman, Dave Levin, Neil Spring
University of Maryland, College Park

Uses of 802.11 packet traces

- MAC Layer (Mahajan et al, Jardosh et al)
- Performance (Rodrig et al)
- Troubleshooting (Cheng et al)

These studies benefit from complete packet traces

What is an incomplete trace?

Transmissions are within range of the monitor
but packets are missing

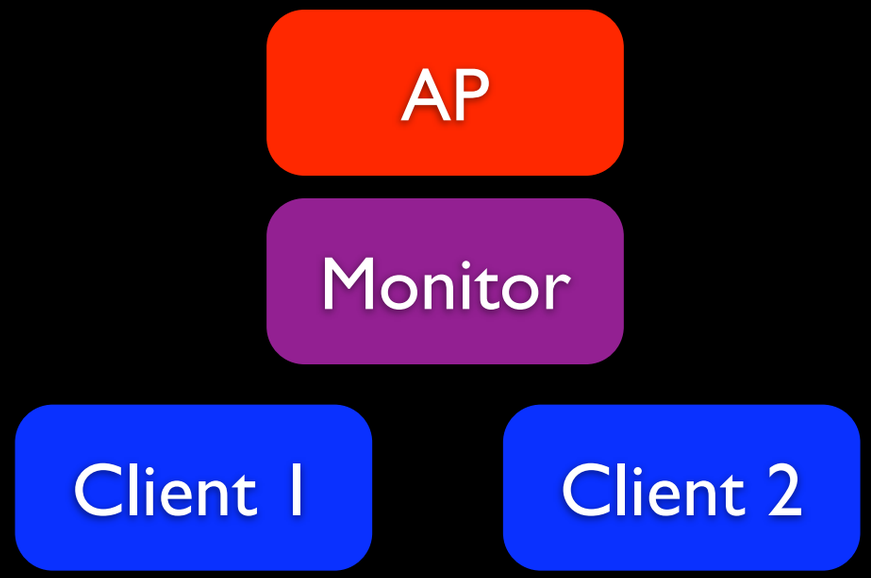

What is an incomplete trace?

Transmissions are within range of the monitor
but packets are missing

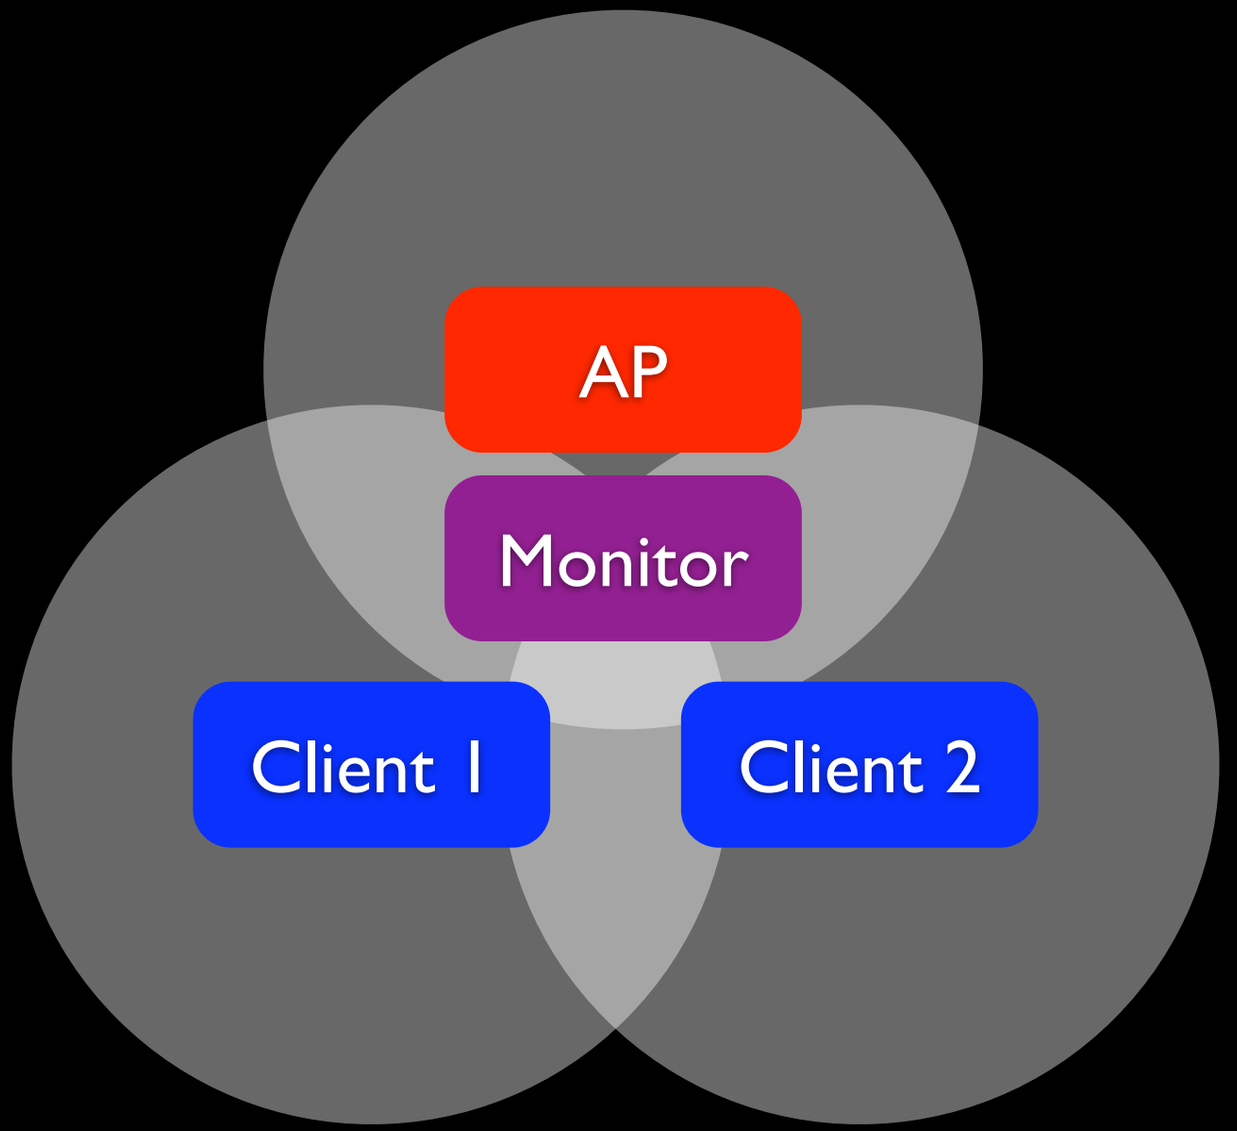

What is an incomplete trace?

Transmissions are within range of the monitor
but packets are missing

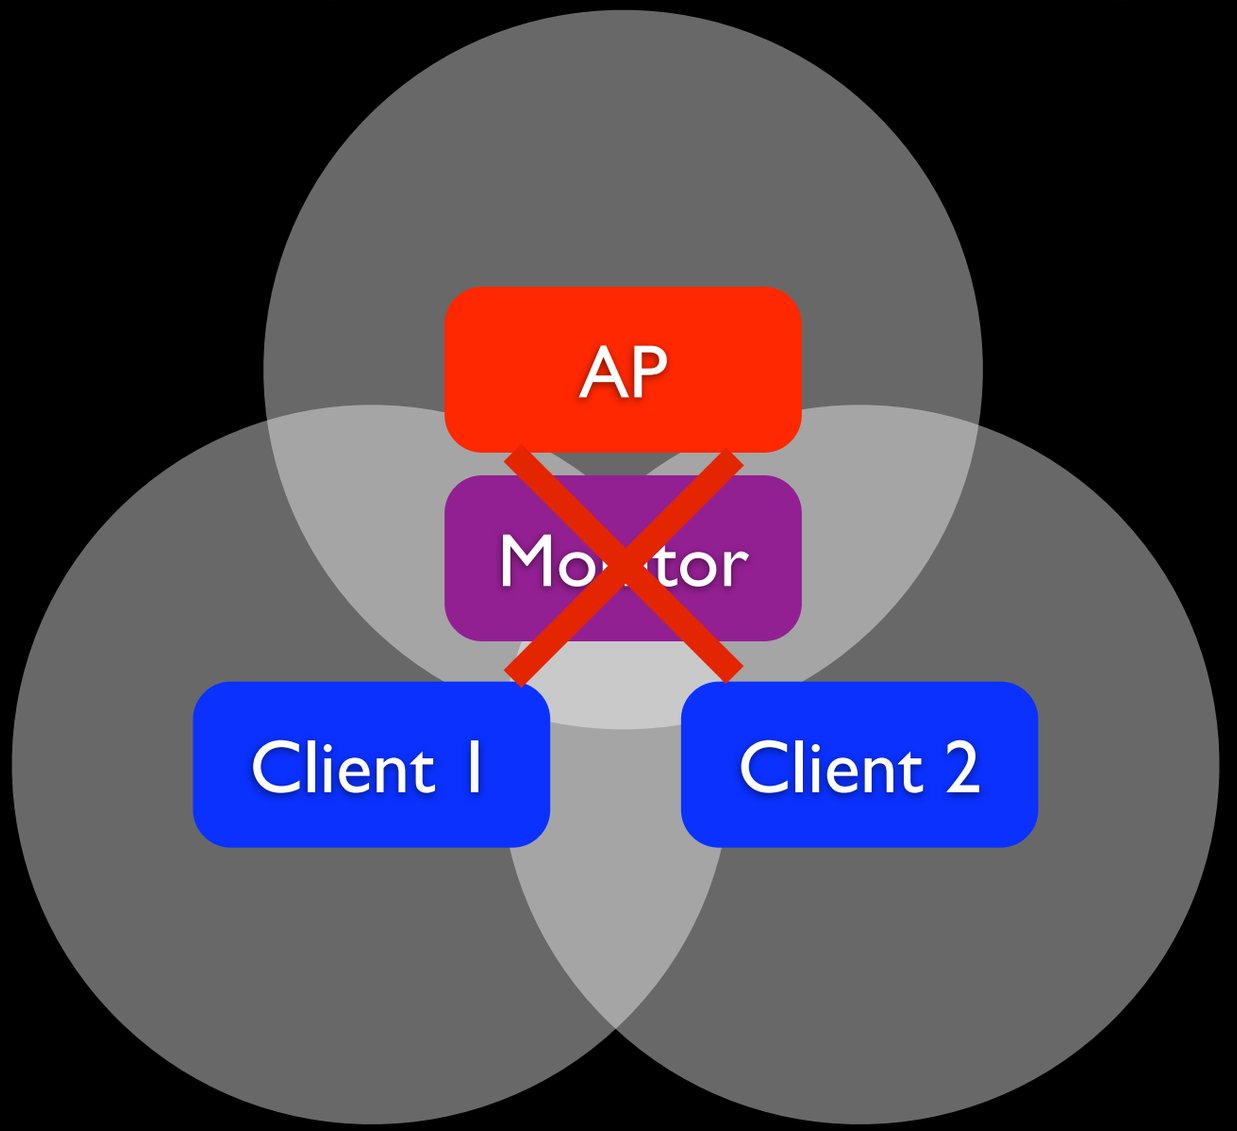

Capturing complete 802.11 packet traces is hard

- Monitor Hardware/Software
- RF Interference
- Monitor Placement
- Merging requires accurate timestamps

(Yeo et al, Portoles-Comeras et al)

Trace Fidelity

Completeness

Did we capture all of the packets?

Accuracy

Did we timestamp the packets correctly?

Main finding: Both are dependent on load

Trace Fidelity

Completeness

Did we capture all of the packets?

Accuracy

Did we timestamp the packets correctly?

Completeness

Did we capture all of
the packets?

Monitors can miss packets

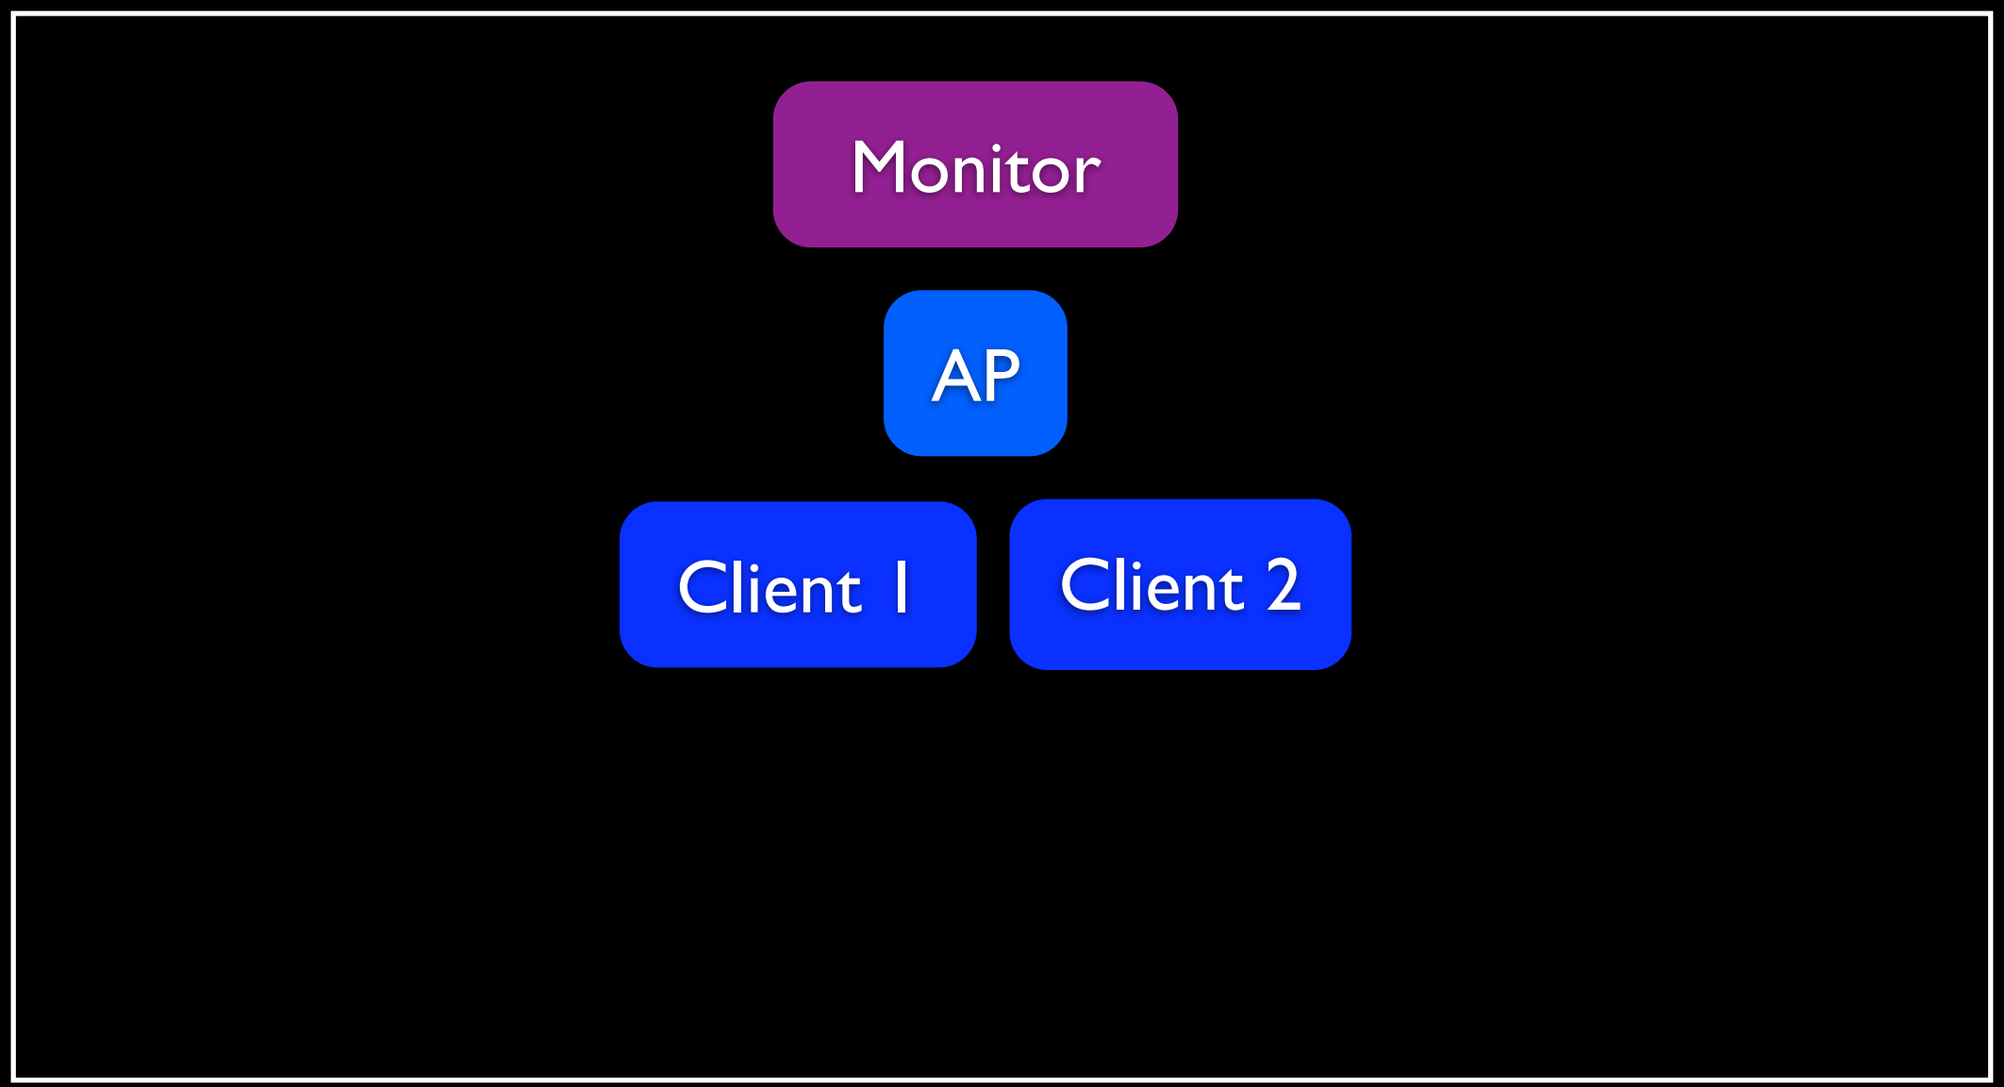

Monitors can miss packets

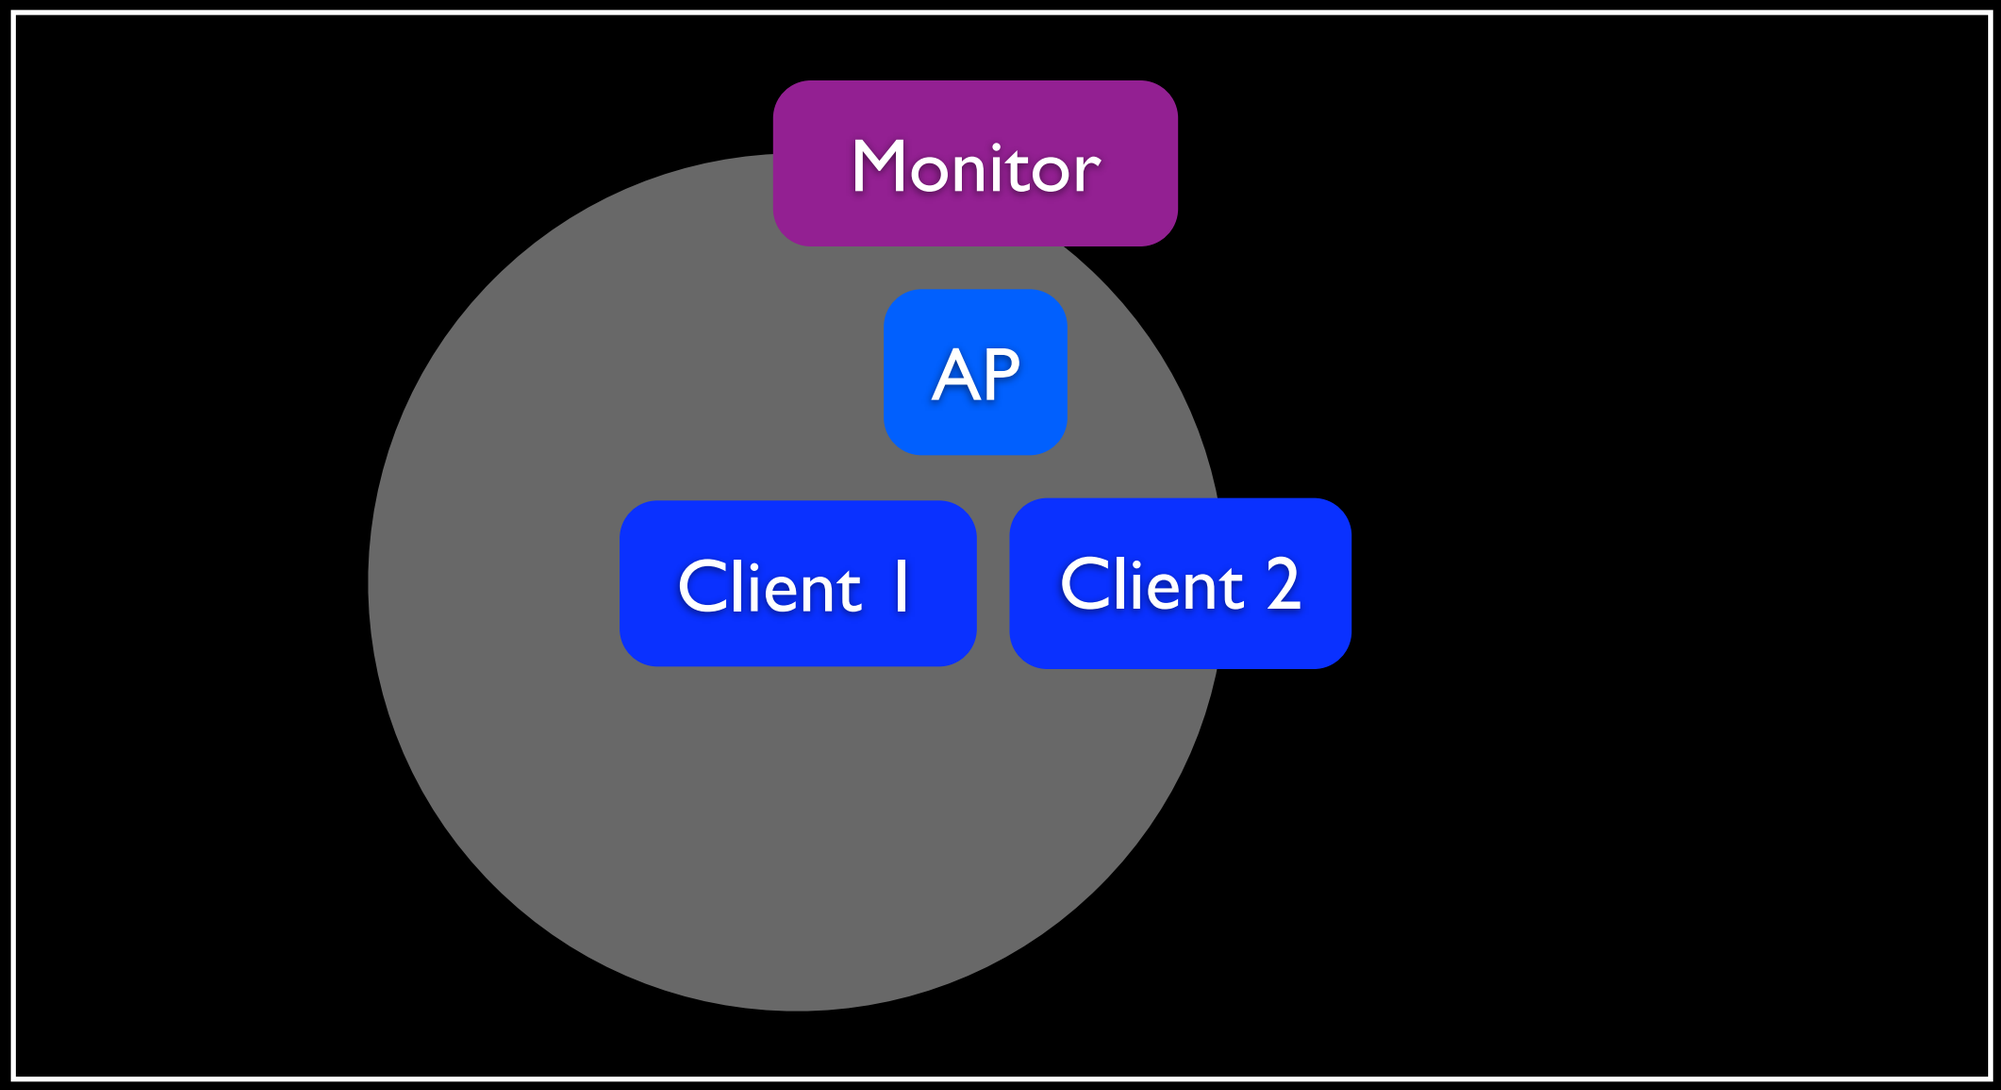

Monitors can miss packets

Both the Monitor and AP receive a packet from Client 1

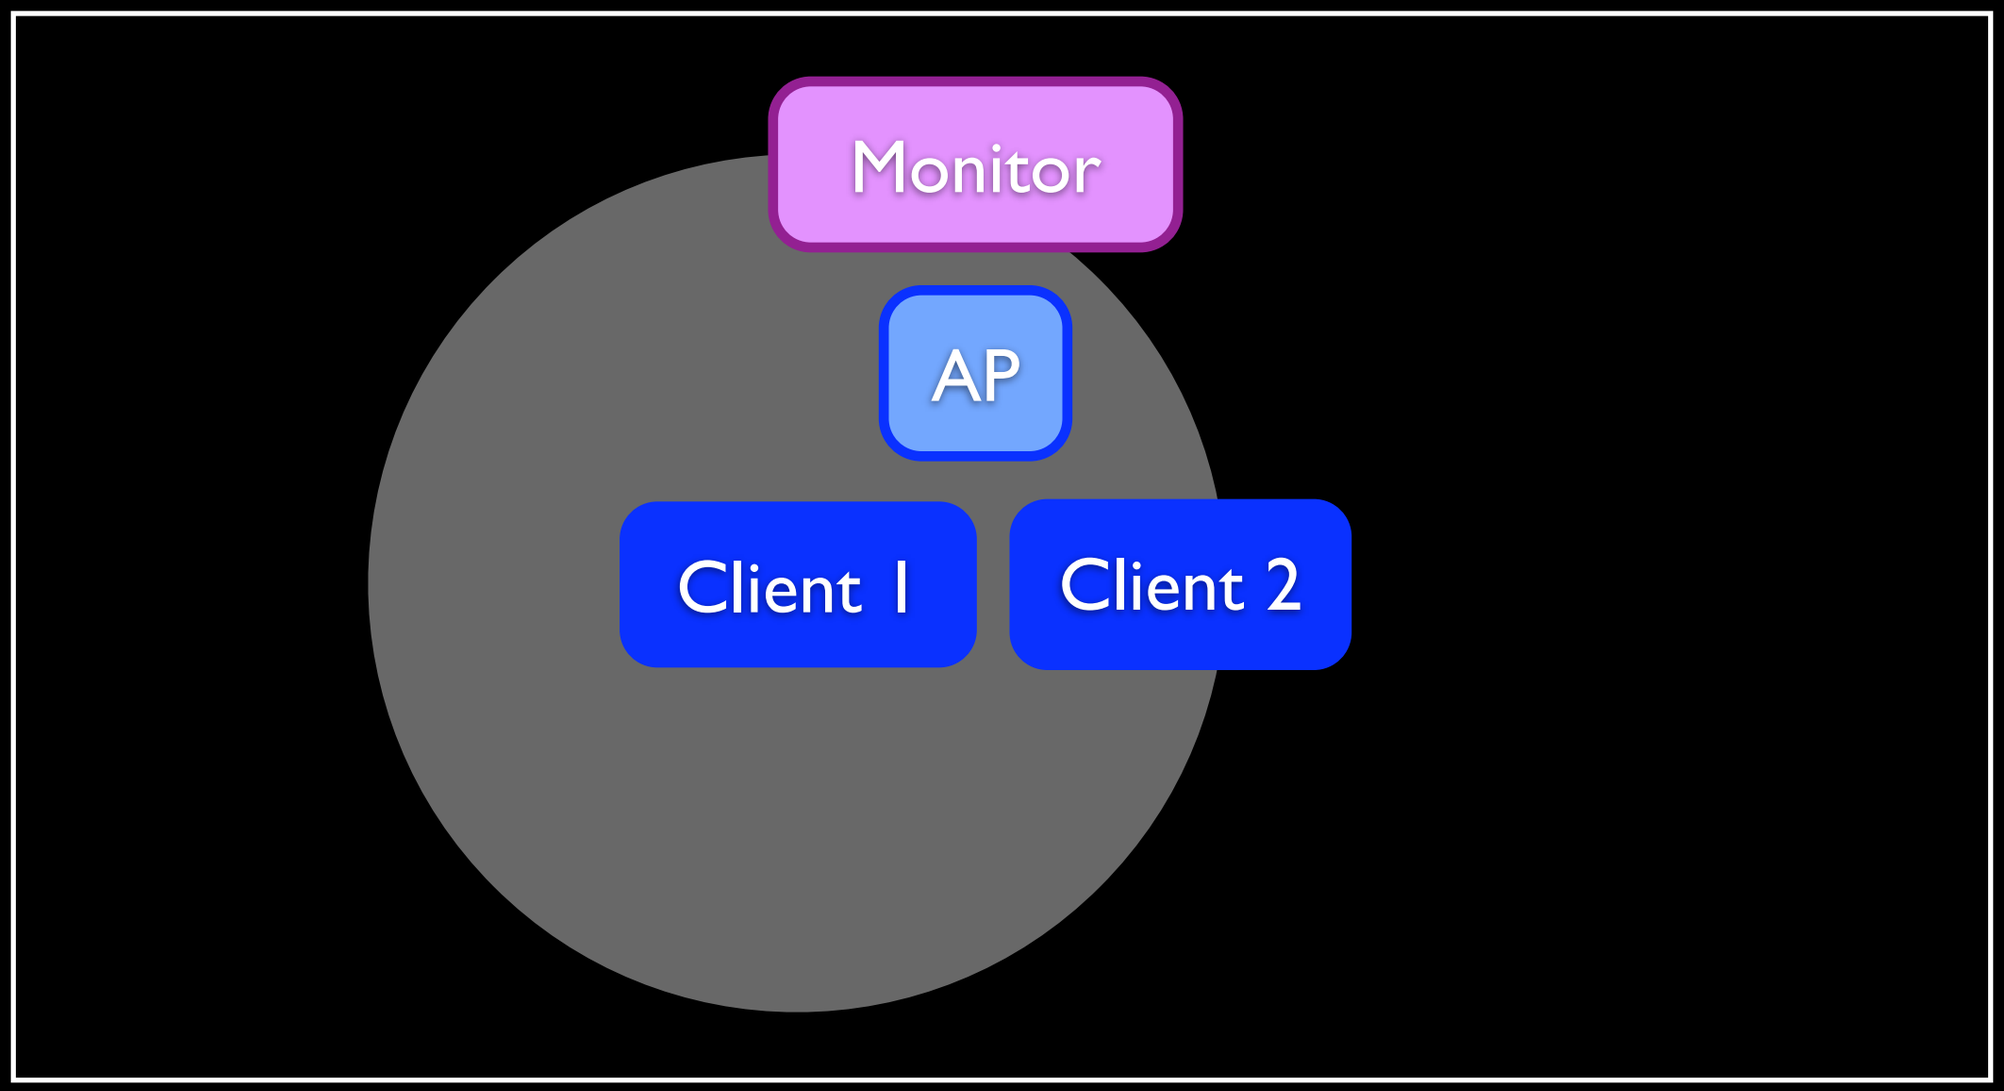

Monitors can miss packets

Both the Monitor and AP receive a packet from Client 1

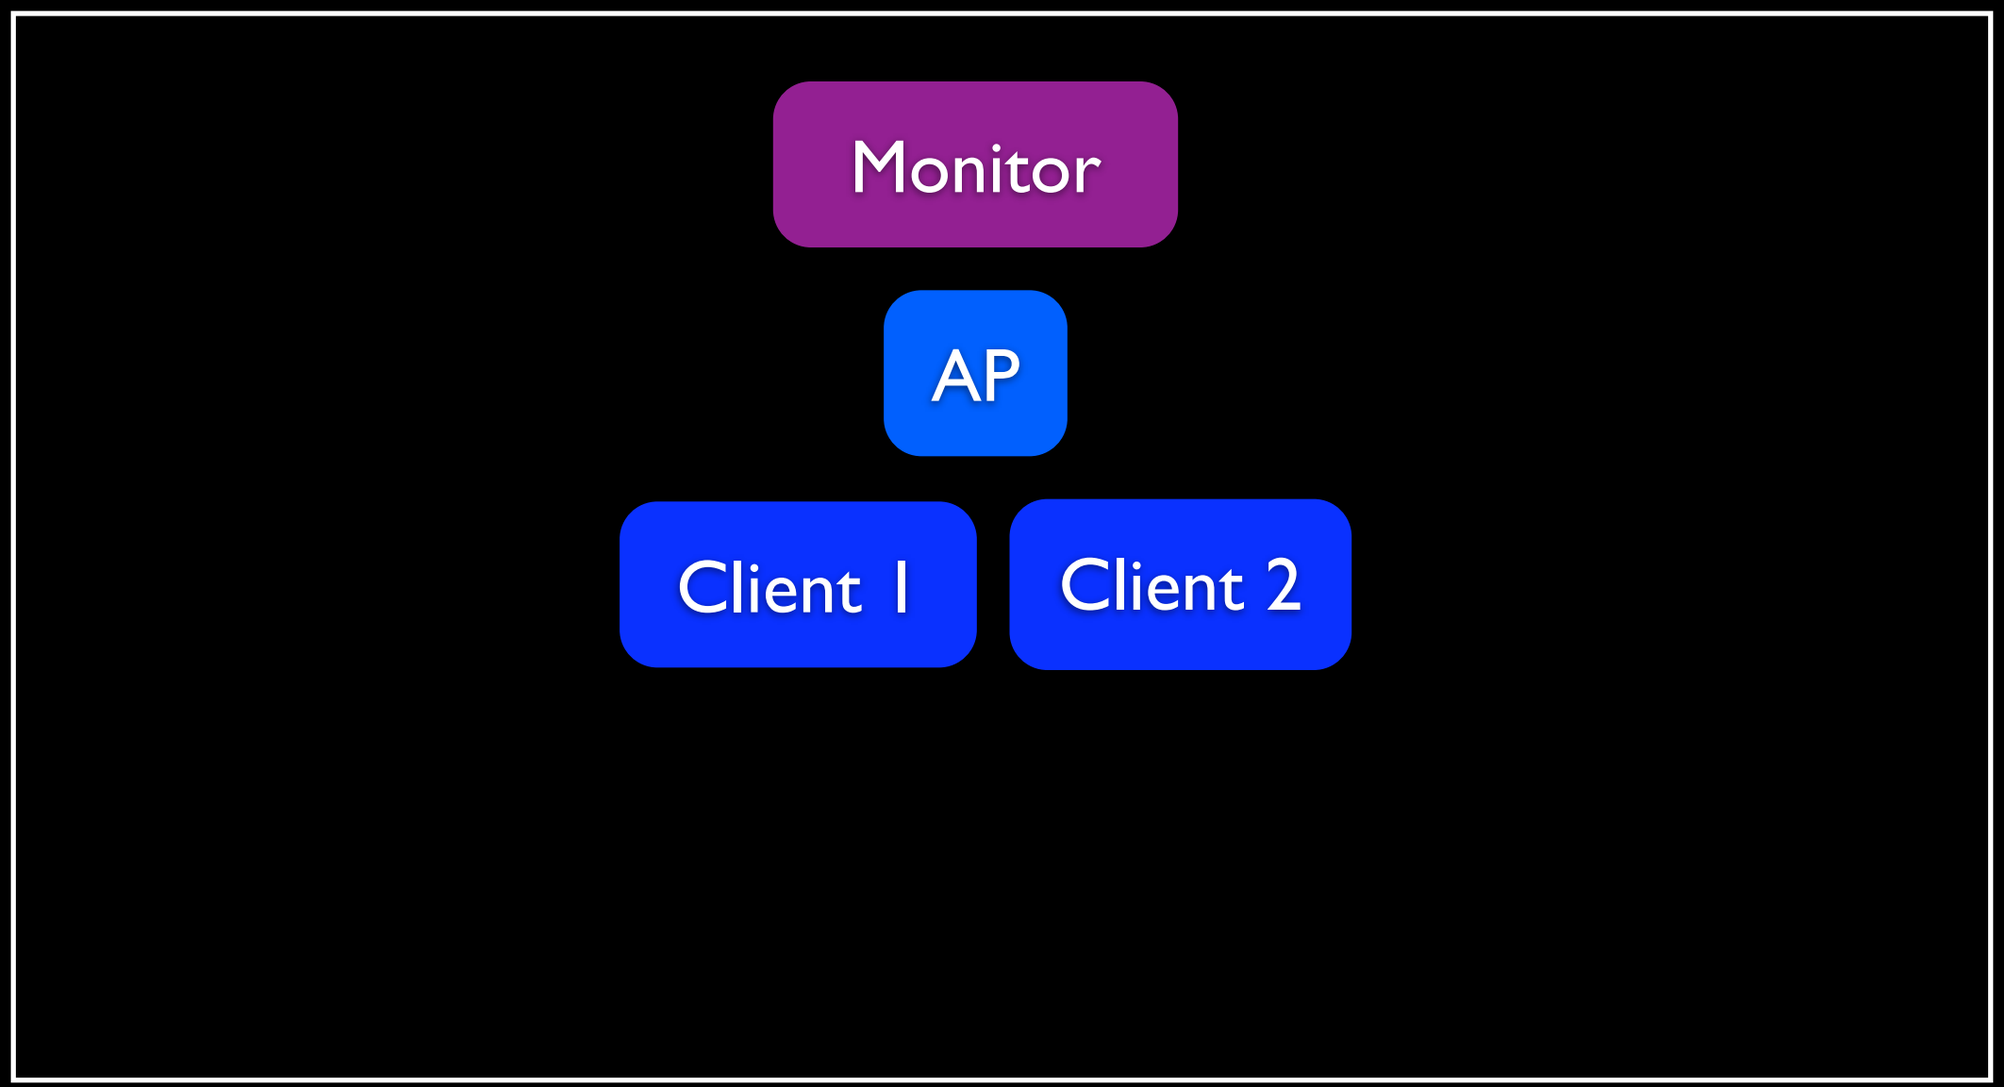

Monitors can miss packets

Both the Monitor and AP receive a packet from Client 1

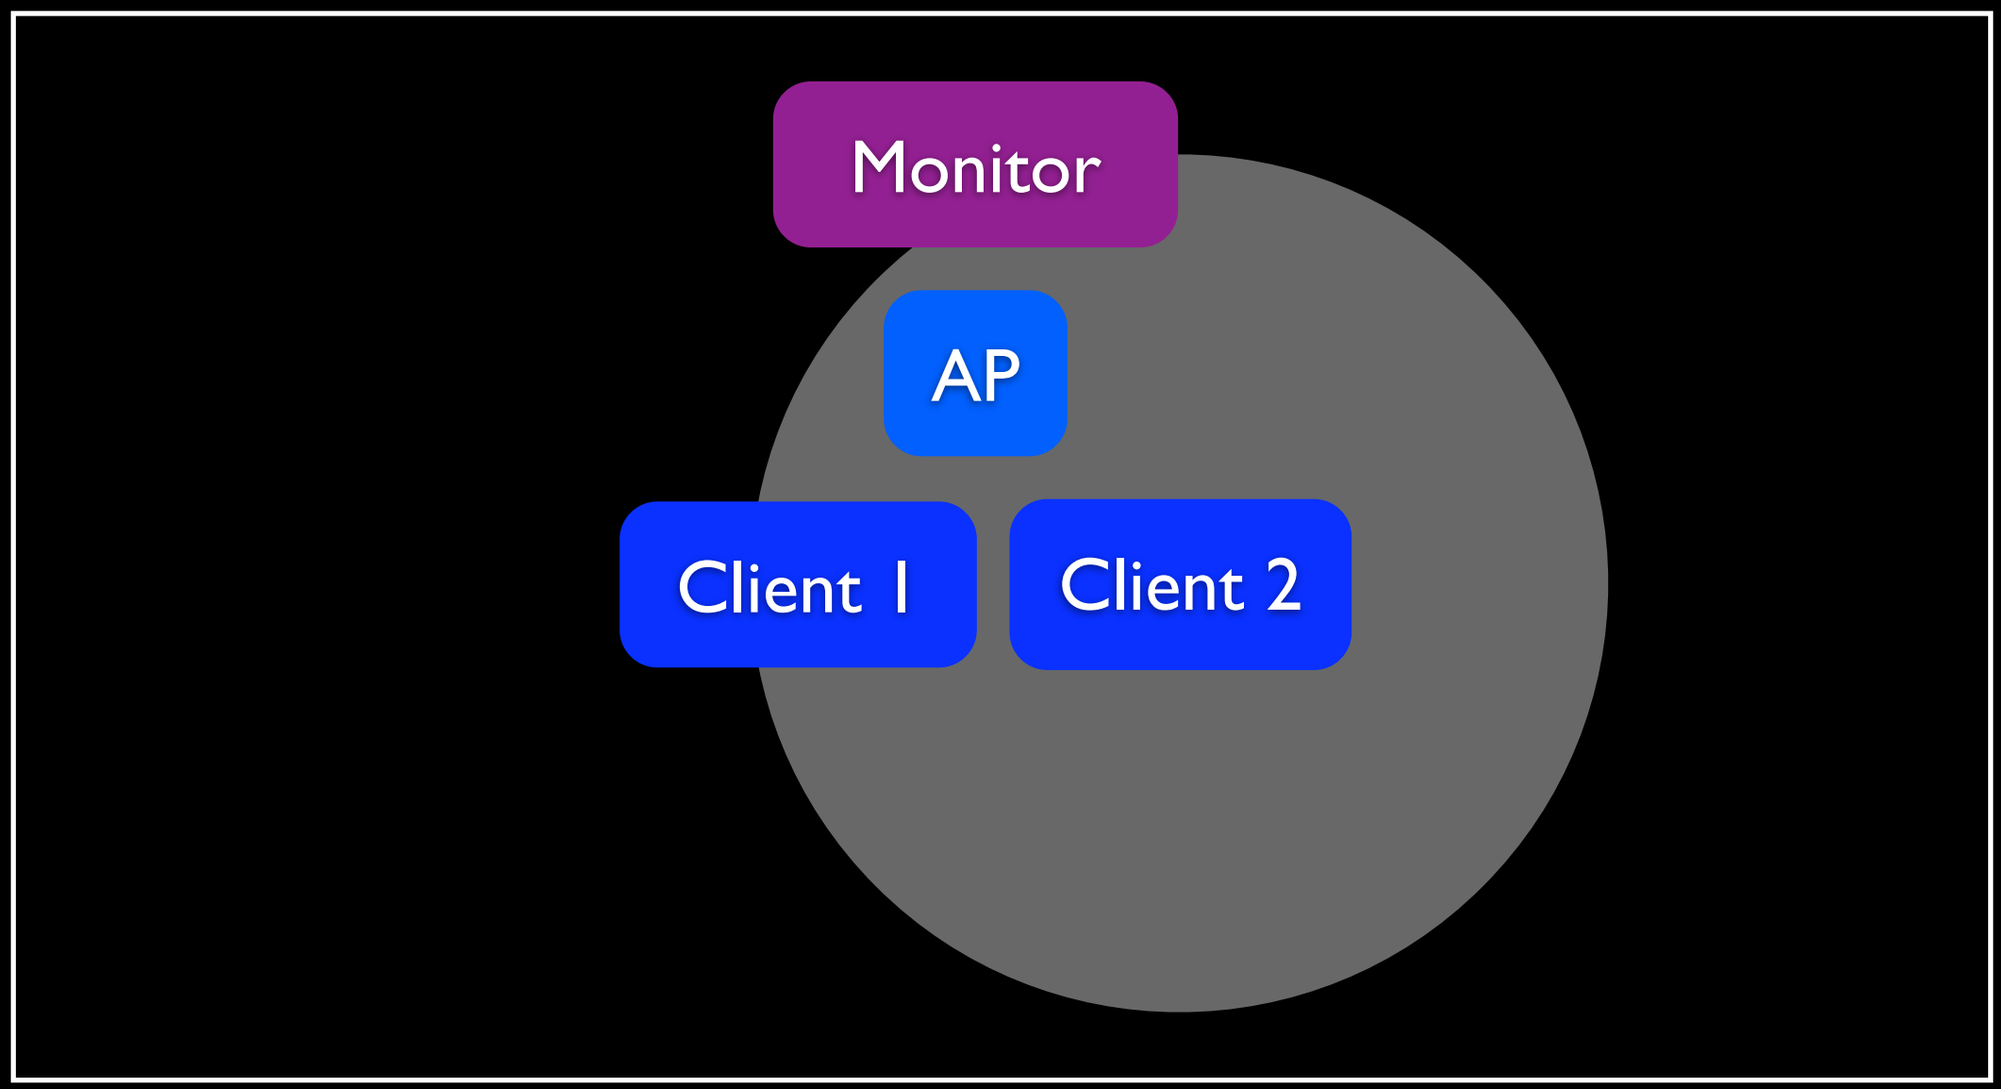

Monitors can miss packets

Both the Monitor and AP receive a packet from Client 1

The Monitor misses a packet from Client 2

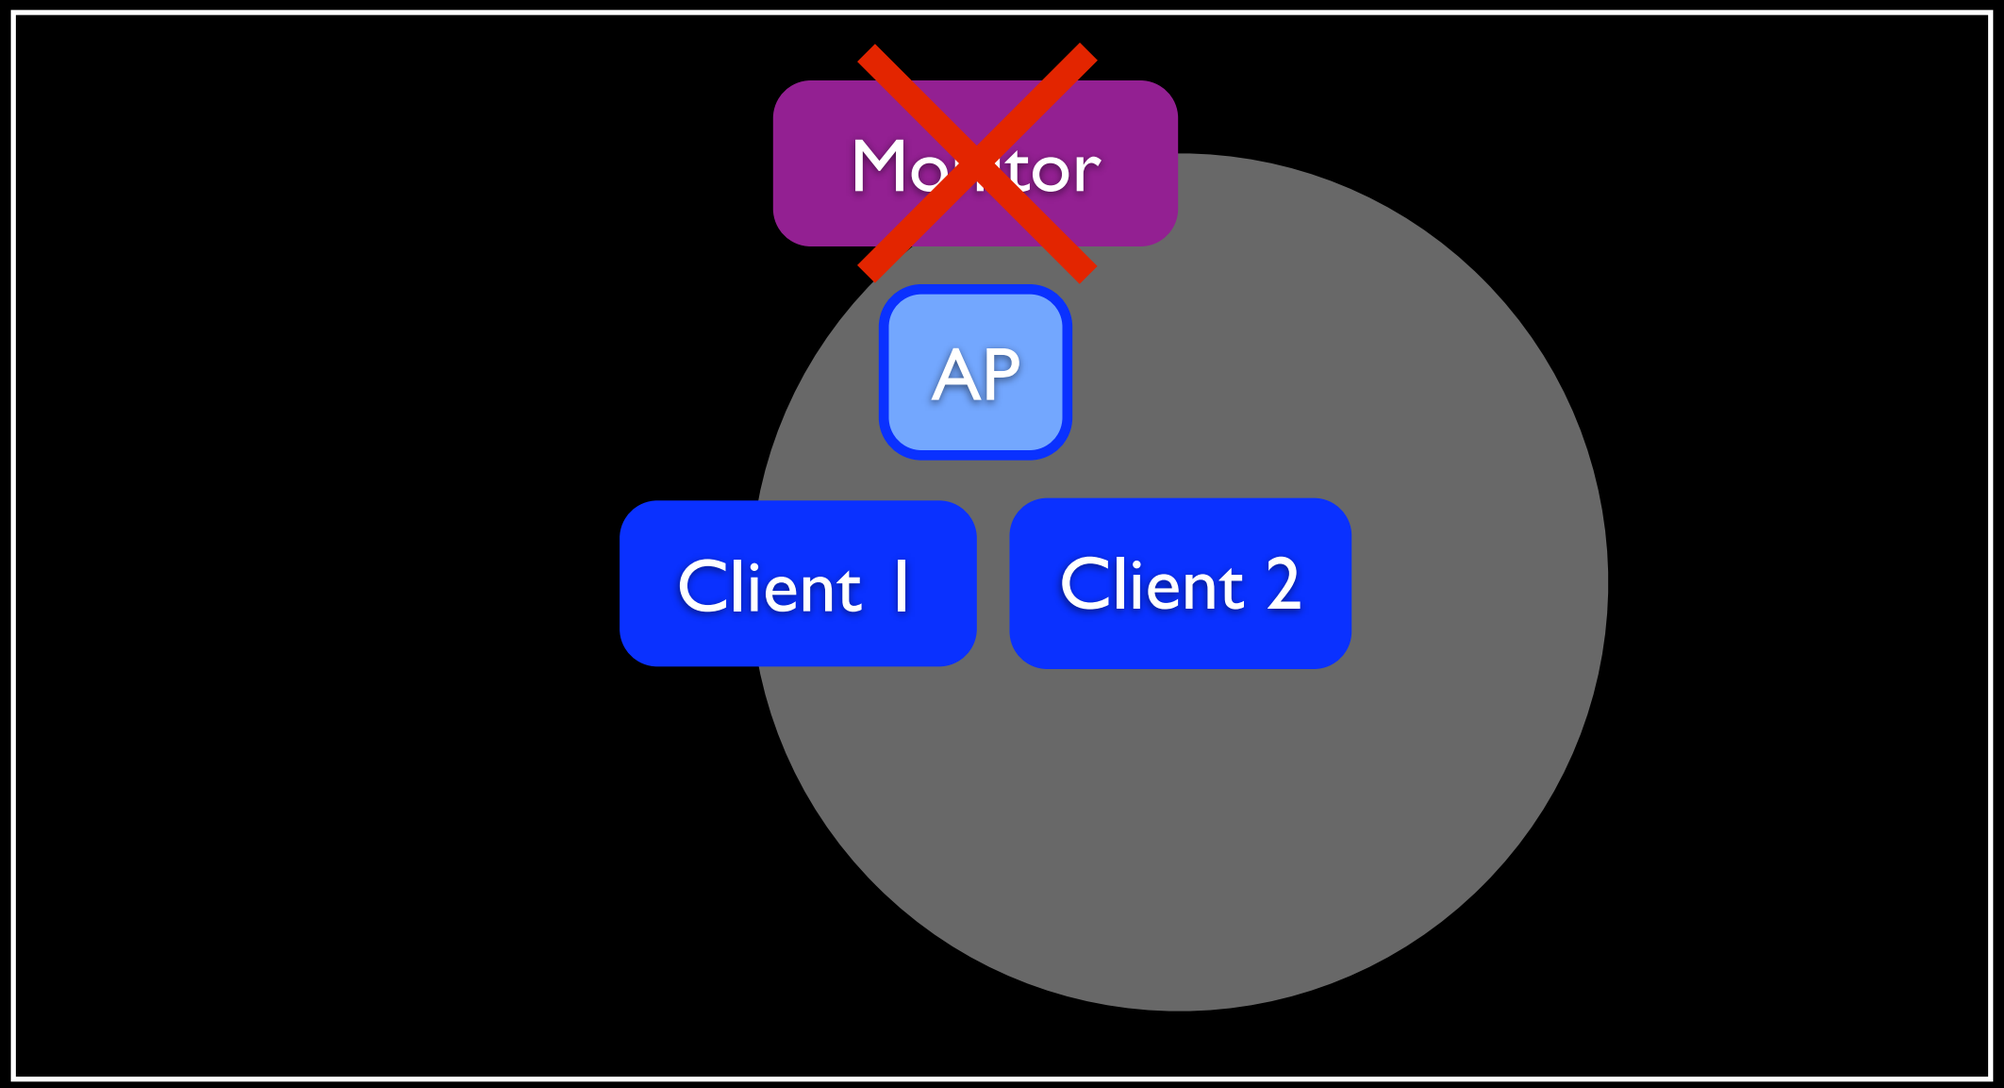

802.11 protocol can show completeness

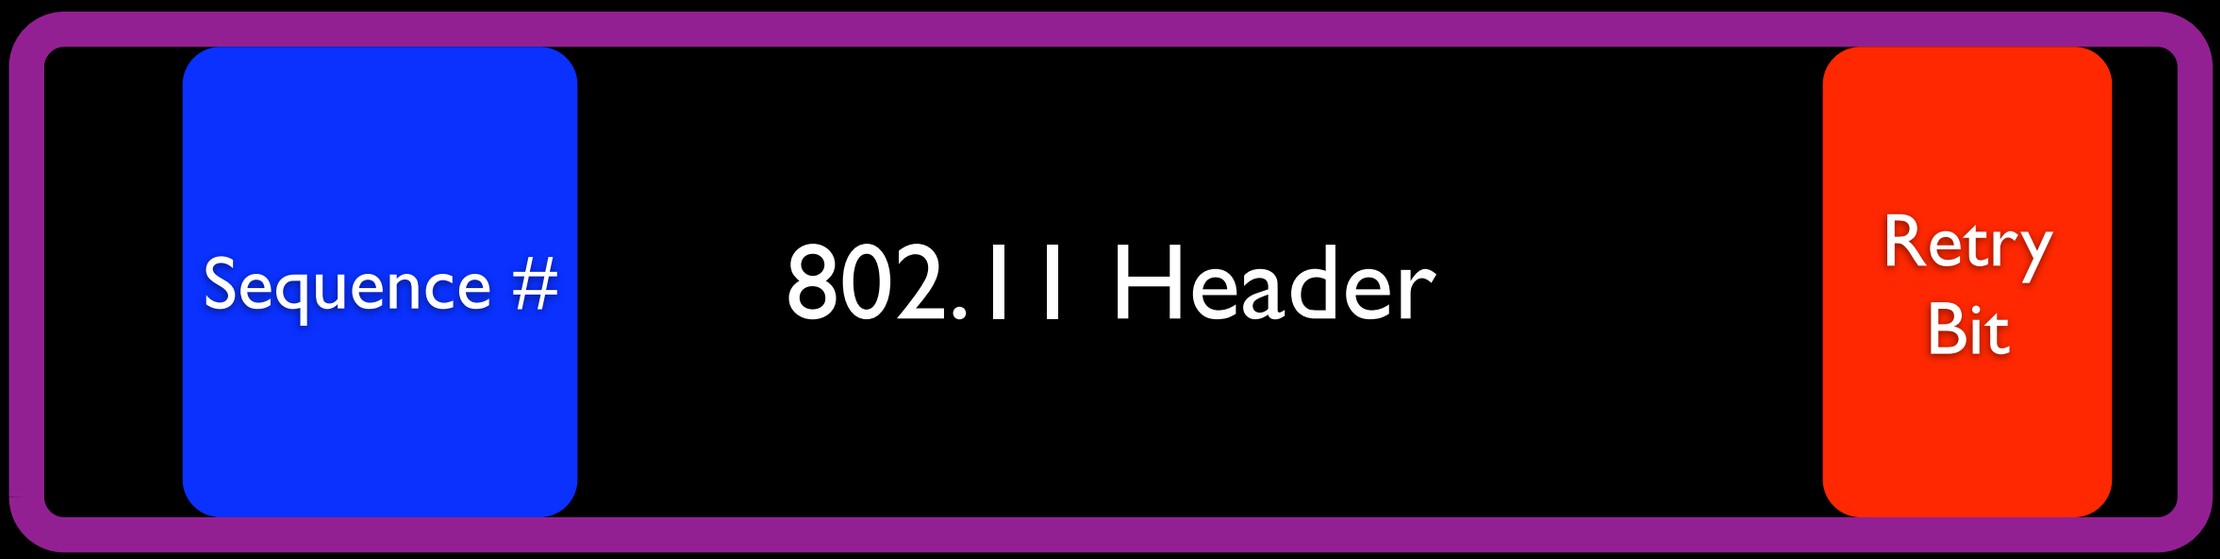

Incremented when a packet is sent

Set when a packet is a retransmission

(Yeo et al)

Estimating completeness

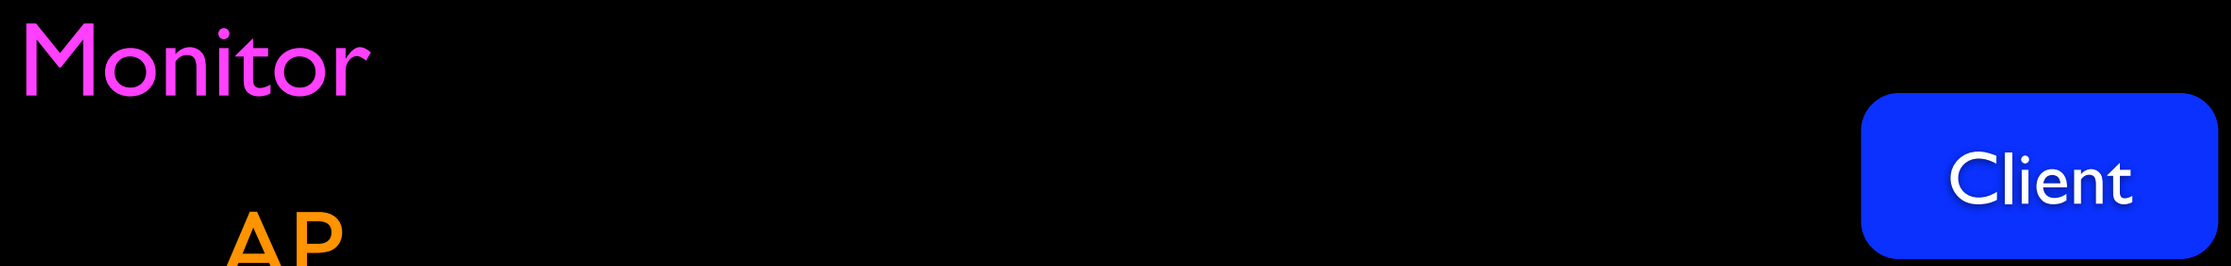

Estimating completeness

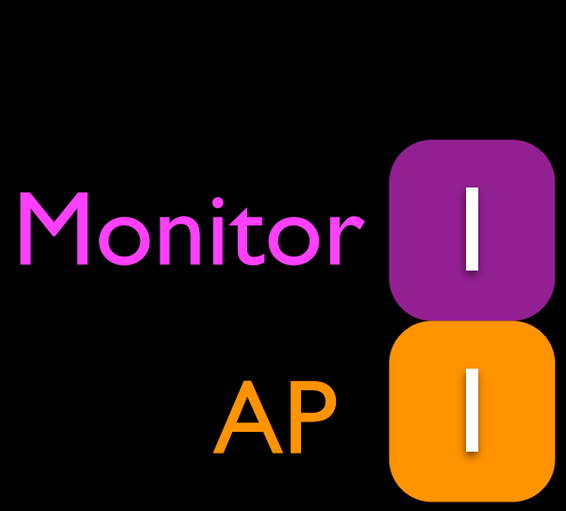

Estimating completeness

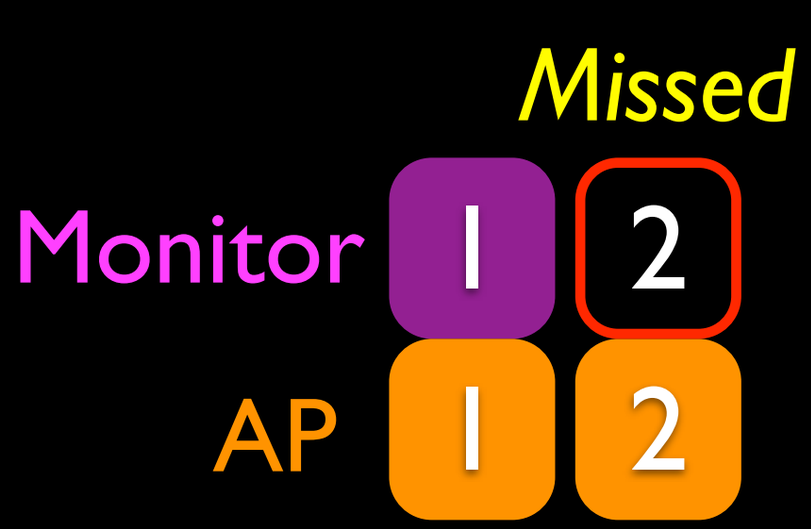

Estimating completeness

2 is missing

Estimating completeness

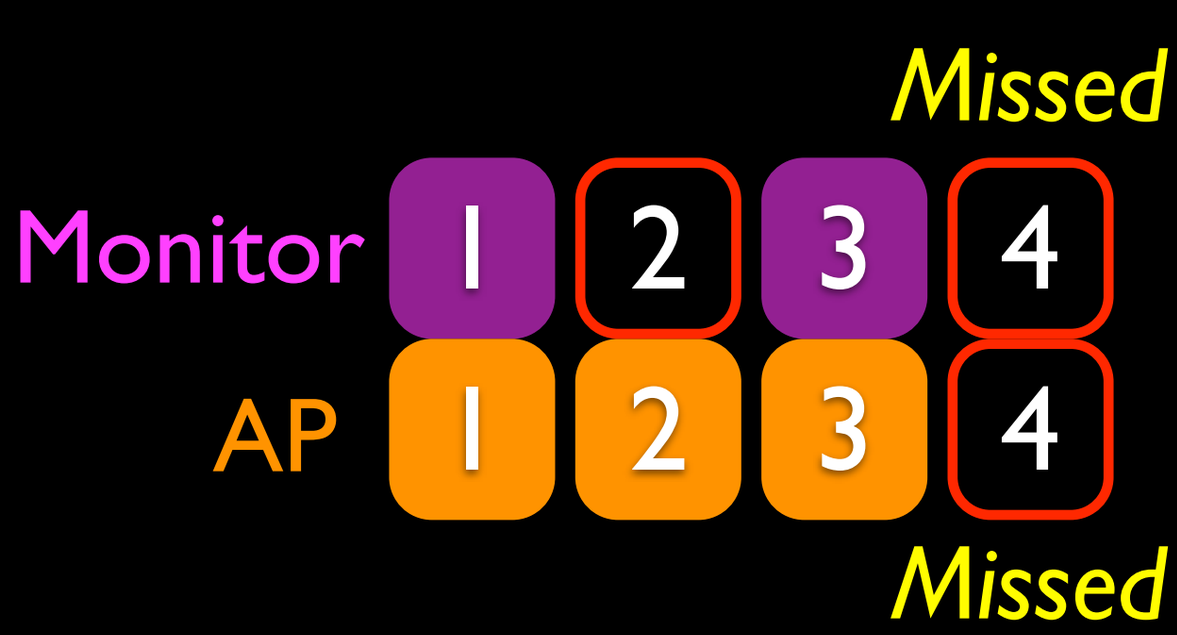

Estimating completeness

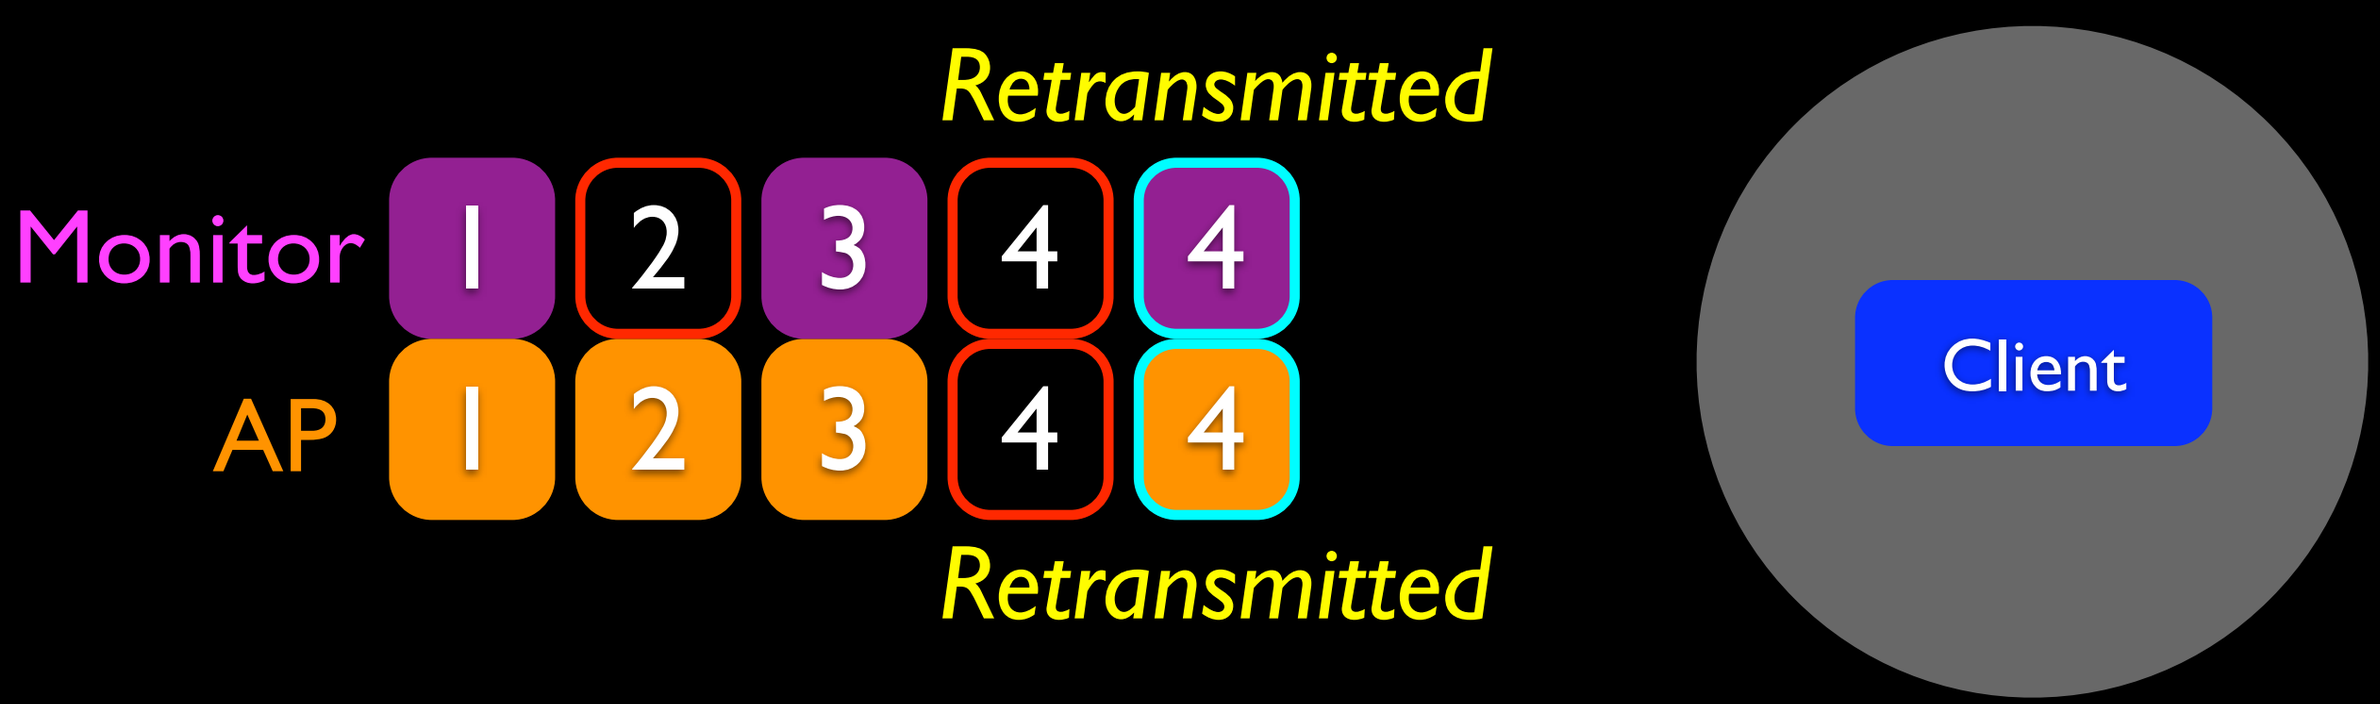

Estimating completeness

4 is missing

Monitor	1	2	3	4	4
AP	1	2	3	4	4

Estimating completeness

The sequence number and retransmission bit show packets 2 and 4 are missing.

How accurate is the estimate?

- Start with SIGCOMM '04 trace CHI
- Randomly removed packets from trace
- Compute estimated # of packets missing
- Relative Error of Method = $\frac{\text{Estimate - Known}}{\text{Total packets}}$

Accuracy of estimate

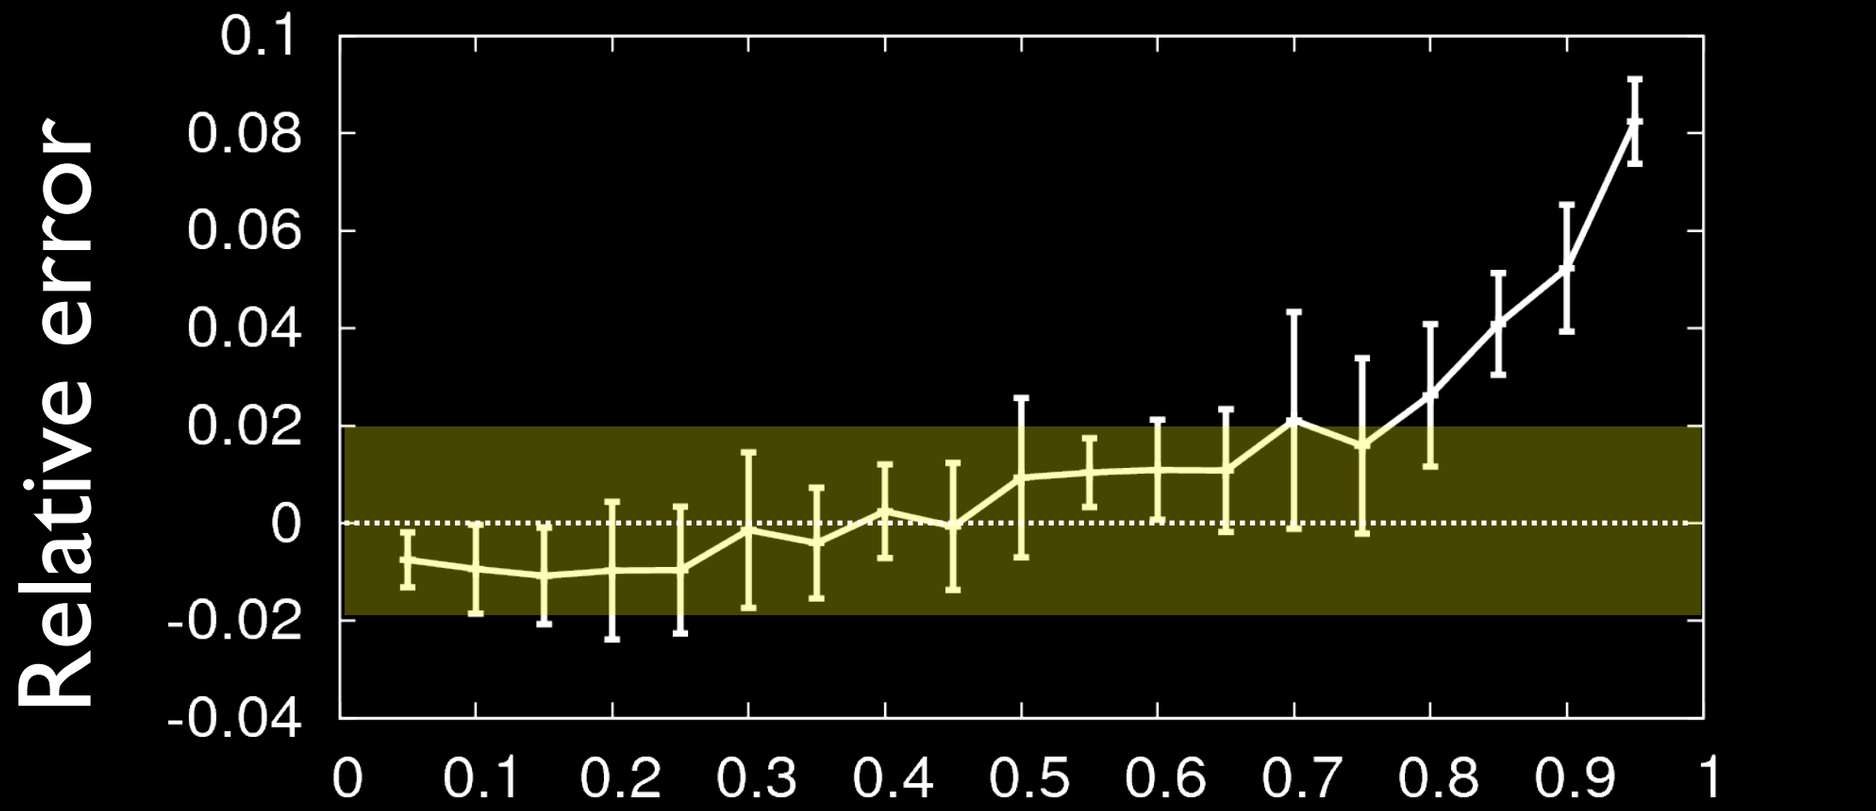

Error bars show 95% confidence interval

Accuracy of estimate

The relative error is < 0.02 when
up to 55% of the trace is removed.

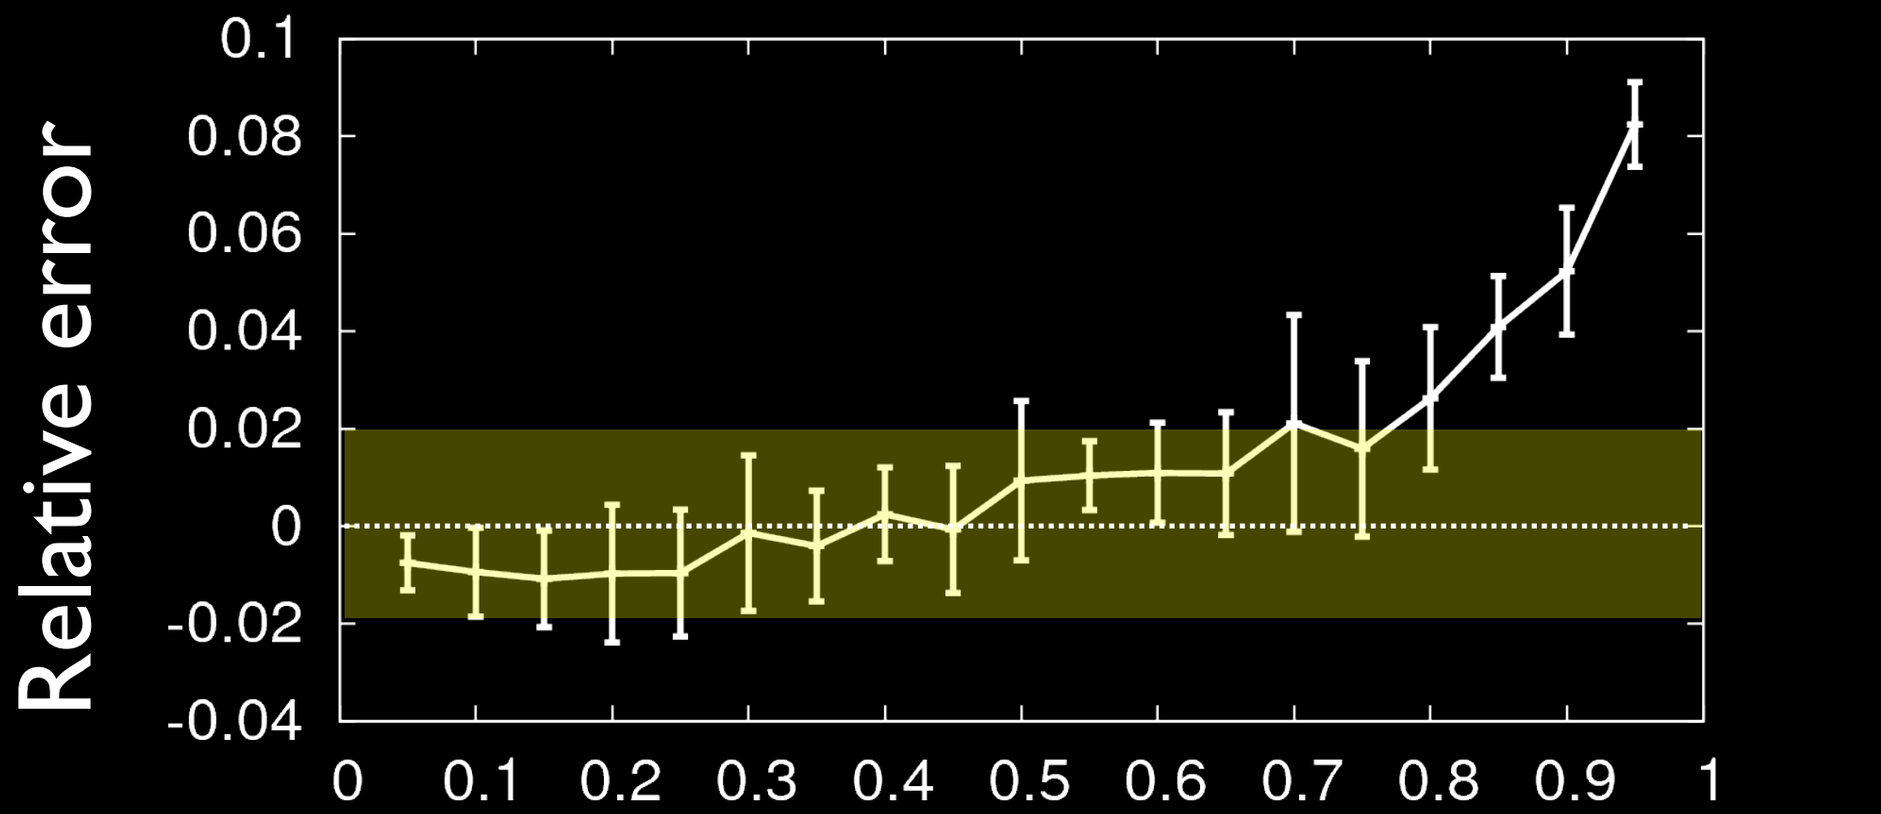

Error bars show 95% confidence interval

Percentage for trace completeness

SIGCOMM 2004 Dataset
Rodrig et al

Percentage for trace completeness

Using the estimate the trace has
of the packets sent by the AP **81%**

SIGCOMM 2004 Dataset
Rodrig et al

Percentage for trace completeness

Using the estimate the trace has
of the packets sent by the AP **81%**

37% of the AP's packets were beacon packets sent when the network was idle

SIGCOMM 2004 Dataset
Rodrig et al

Percentage for trace completeness

Using the estimate the trace has
of the packets sent by the AP **81%**

37% of the AP's packets were beacon packets sent when the network was idle

Excluding idle beacon packets
of packets sent by the AP are in the trace **70%**

SIGCOMM 2004 Dataset
Rodrig et al

One number is not enough

- Problem: Completeness is only interesting when the network is under load
  - Example: Capturing a trace from an AP overnight
- Solution: Estimate completeness within small trace intervals
  - Beacons are sent by AP every 100ms

Trace completeness score

$$\frac{\text{Packets collected}_i}{\text{Packets expected}_i}$$

Trace completeness score

For all devices in-range

$$\frac{\text{Packets collected}_i}{\text{Packets expected}_i}$$

Trace completeness score

For all devices in-range

$$\frac{Packets_i}{Packets\ expected_i}$$

Trace completeness score

For all devices in-range

$$\frac{Packets_i}{Sequence\ Change_i + Retransmissions_i}$$

Trace completeness score

For all devices in-range

$$\frac{Packets_i}{Sequence\ Change_i + Retransmissions_i}$$

Quantifies the completeness of interval i

Visualizing trace completeness

Visualizing trace completeness

- Y-Axis: Score
  - Completeness of an Interval

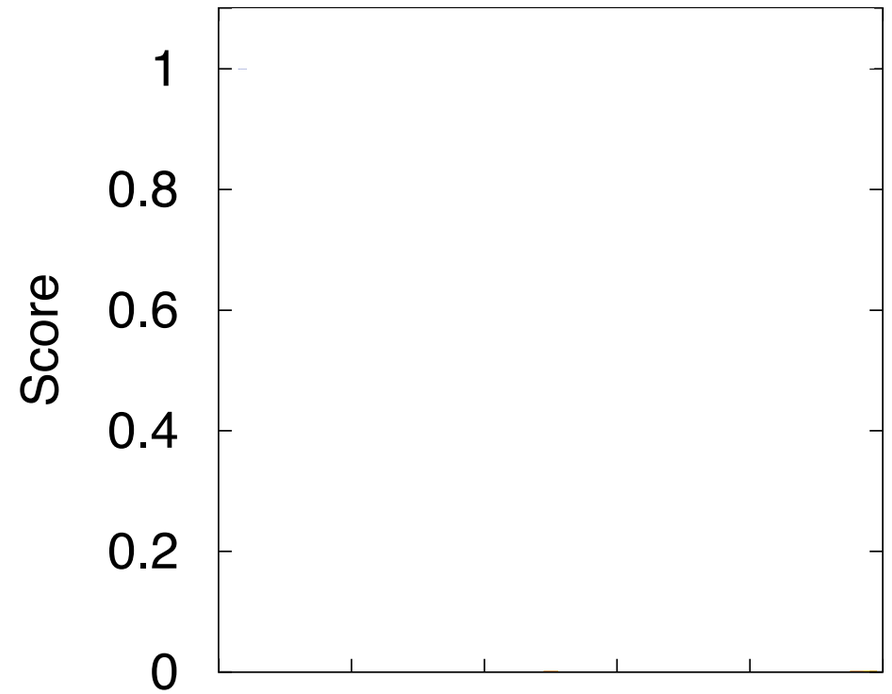

Visualizing trace completeness

- Y-Axis: Score
  - Completeness of an Interval
- X-Axis: Load
  - Sequence # change

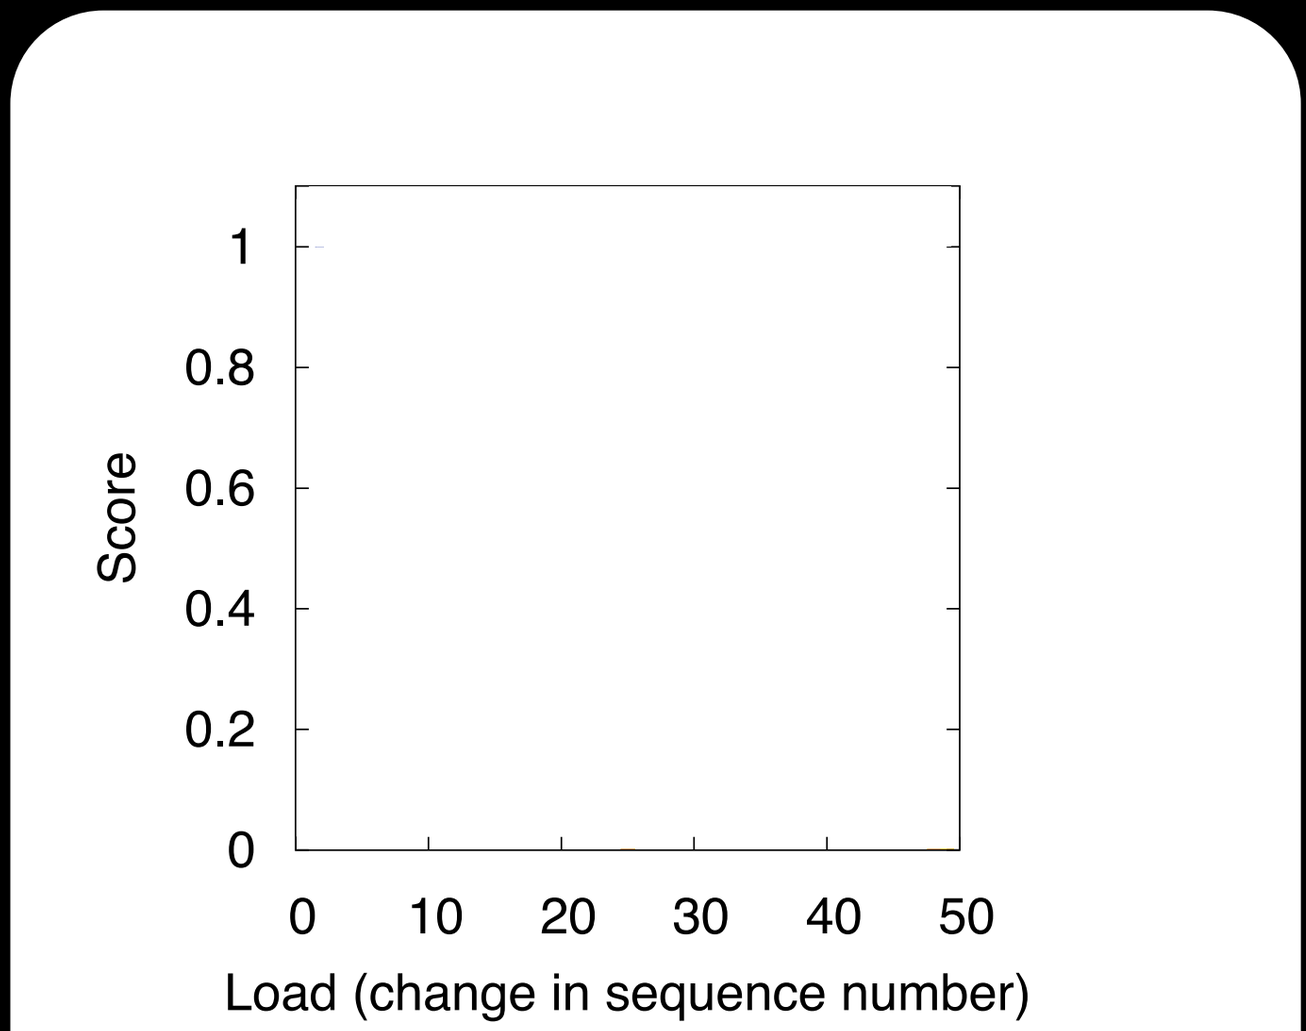

Visualizing trace completeness

- Y-Axis: Score
  - Completeness of an Interval
- X-Axis: Load
  - Sequence # change
- Color: Frequency

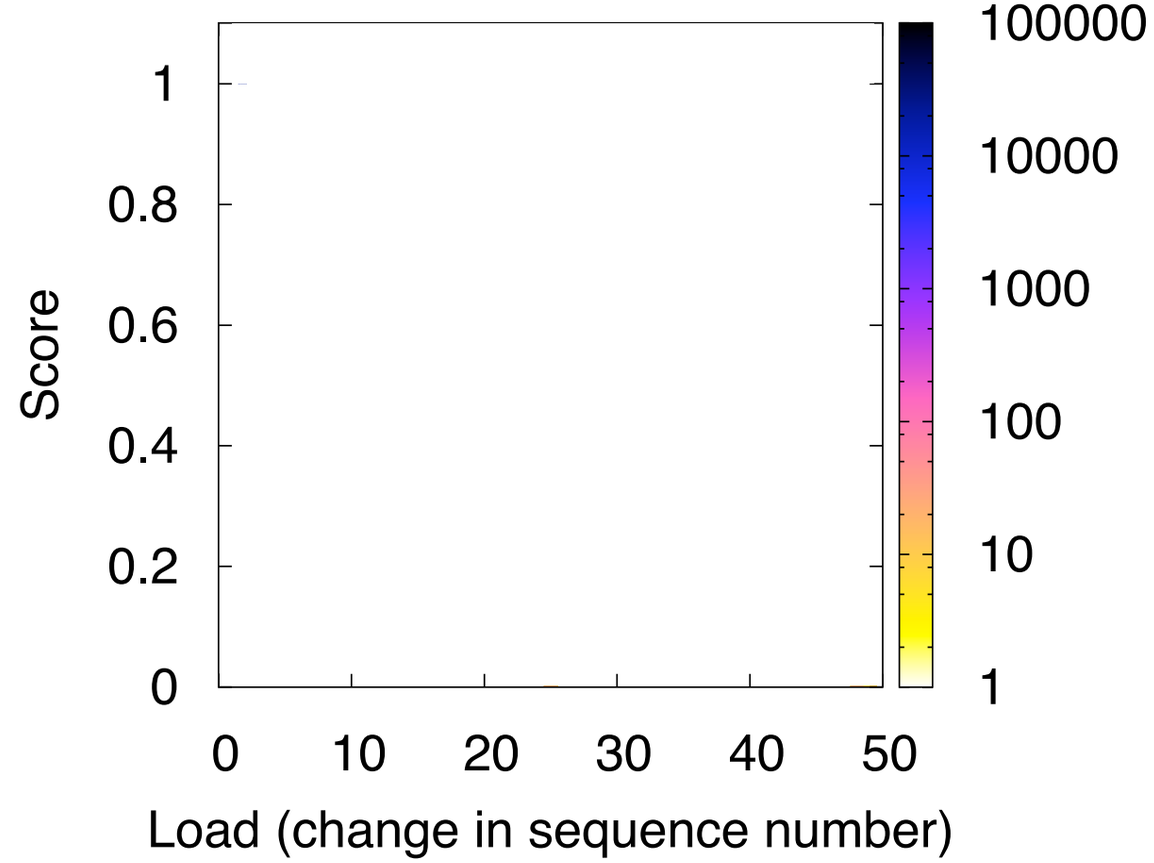

Completeness with T-Fi plot

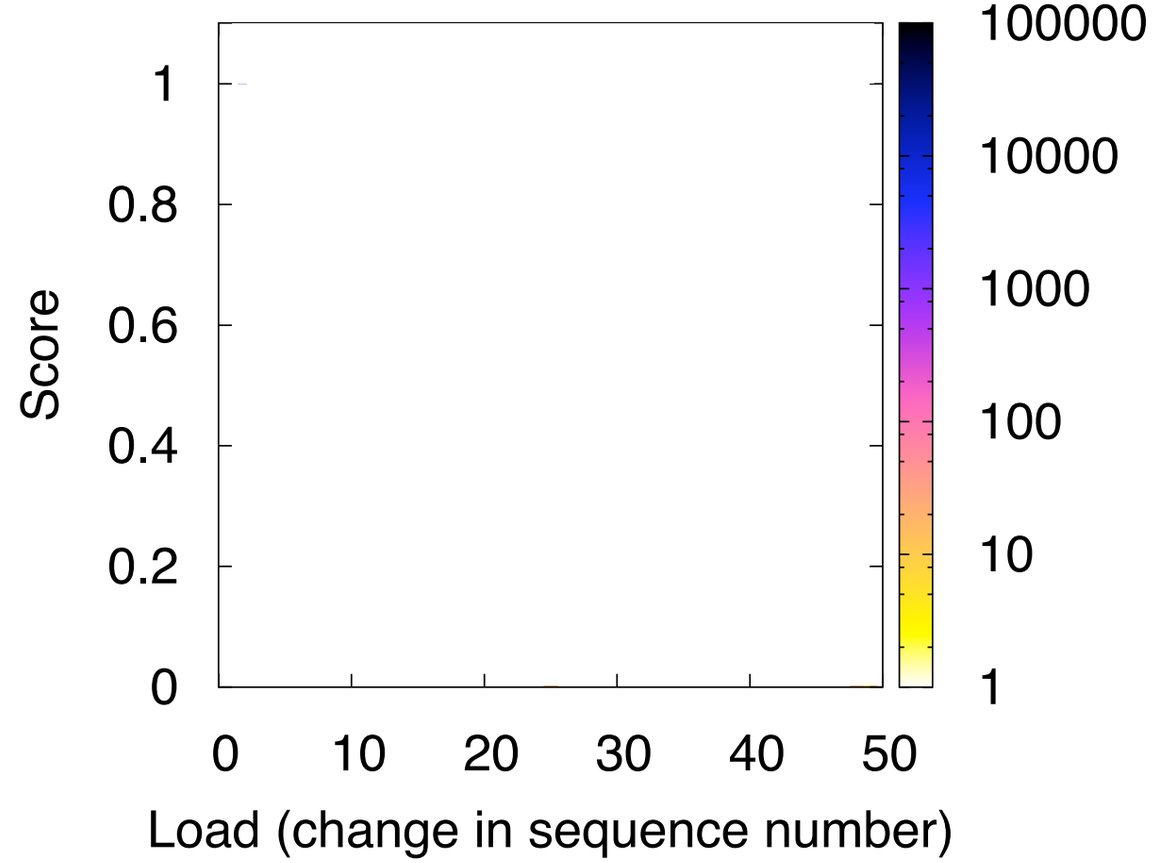

Completeness with T-Fi plot

- Complete loaded trace has dark area on top

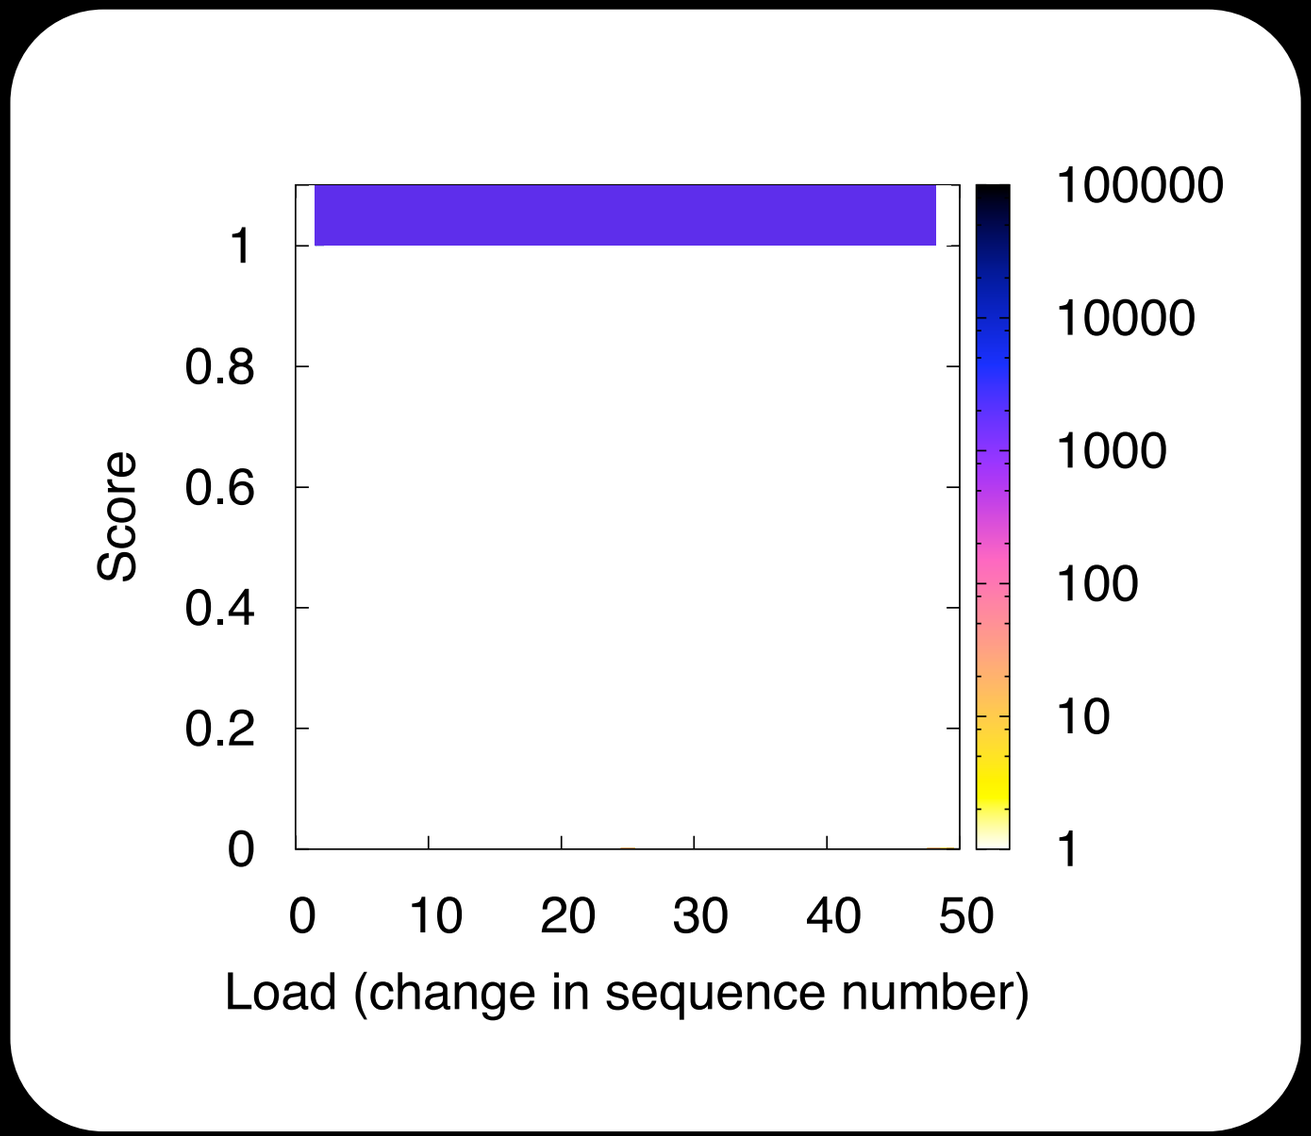

Completeness with T-Fi plot

- Complete loaded trace has dark area on top
- Incomplete trace has lower dark areas

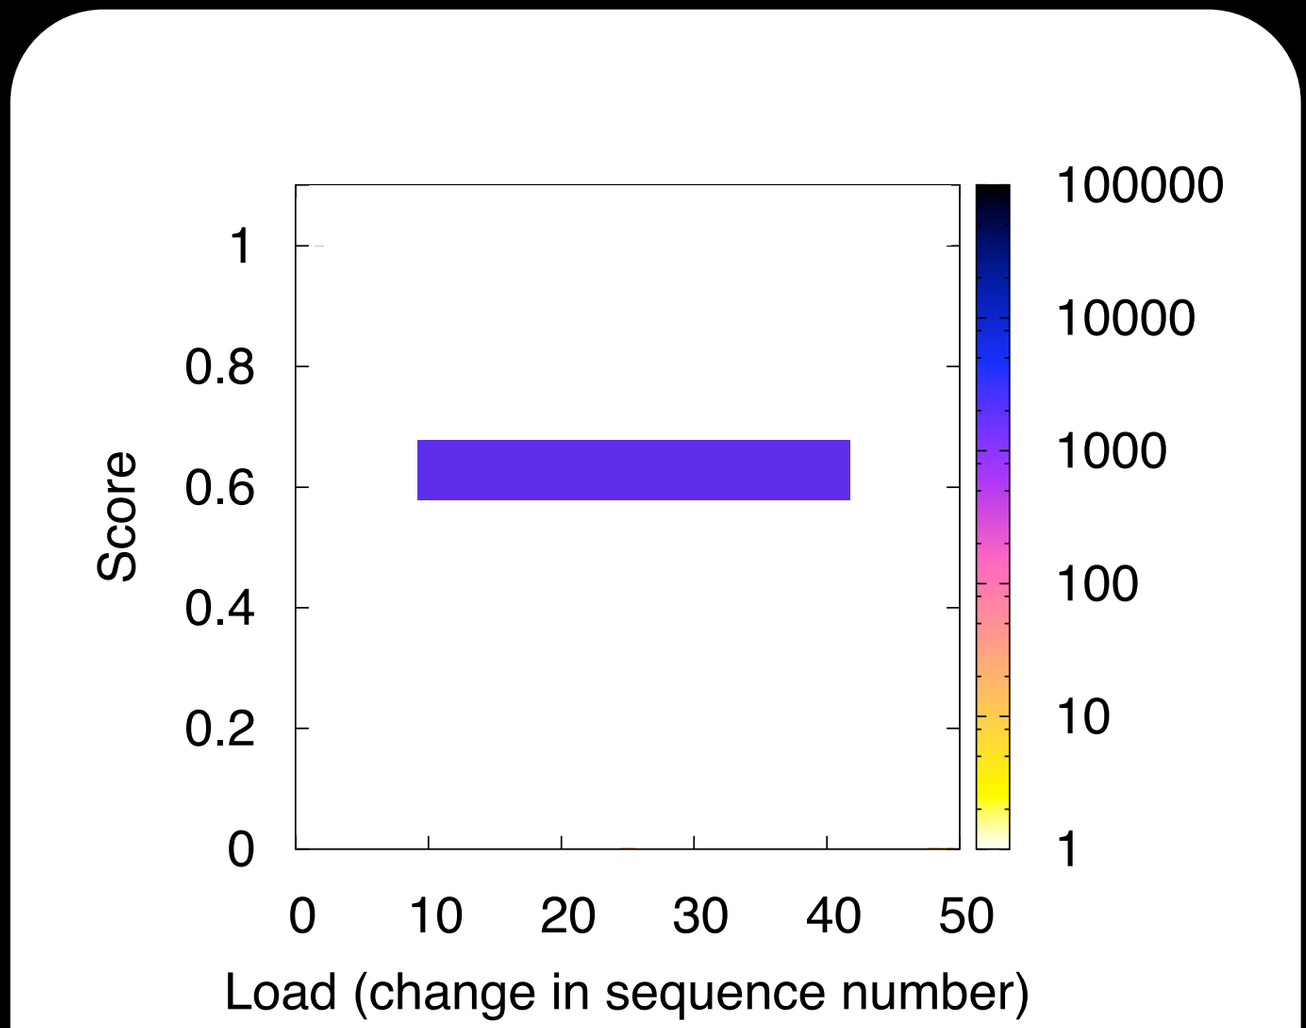

Completeness with T-Fi plot

- Complete loaded trace has dark area on top
- Incomplete trace has lower dark areas
- Low load trace does not have dark color on right

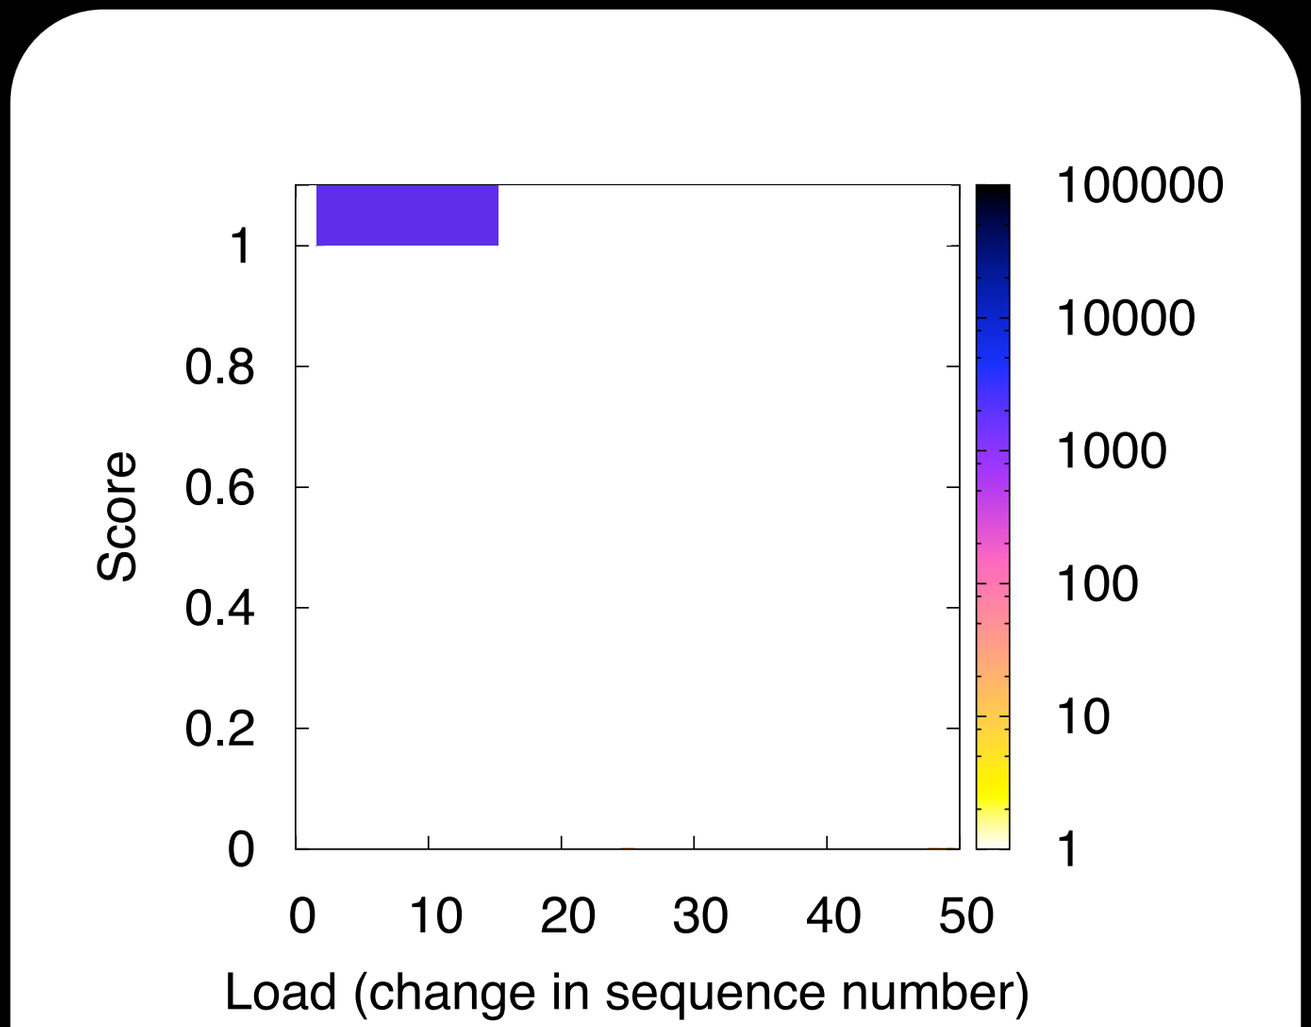

T-Fi plots focus on load

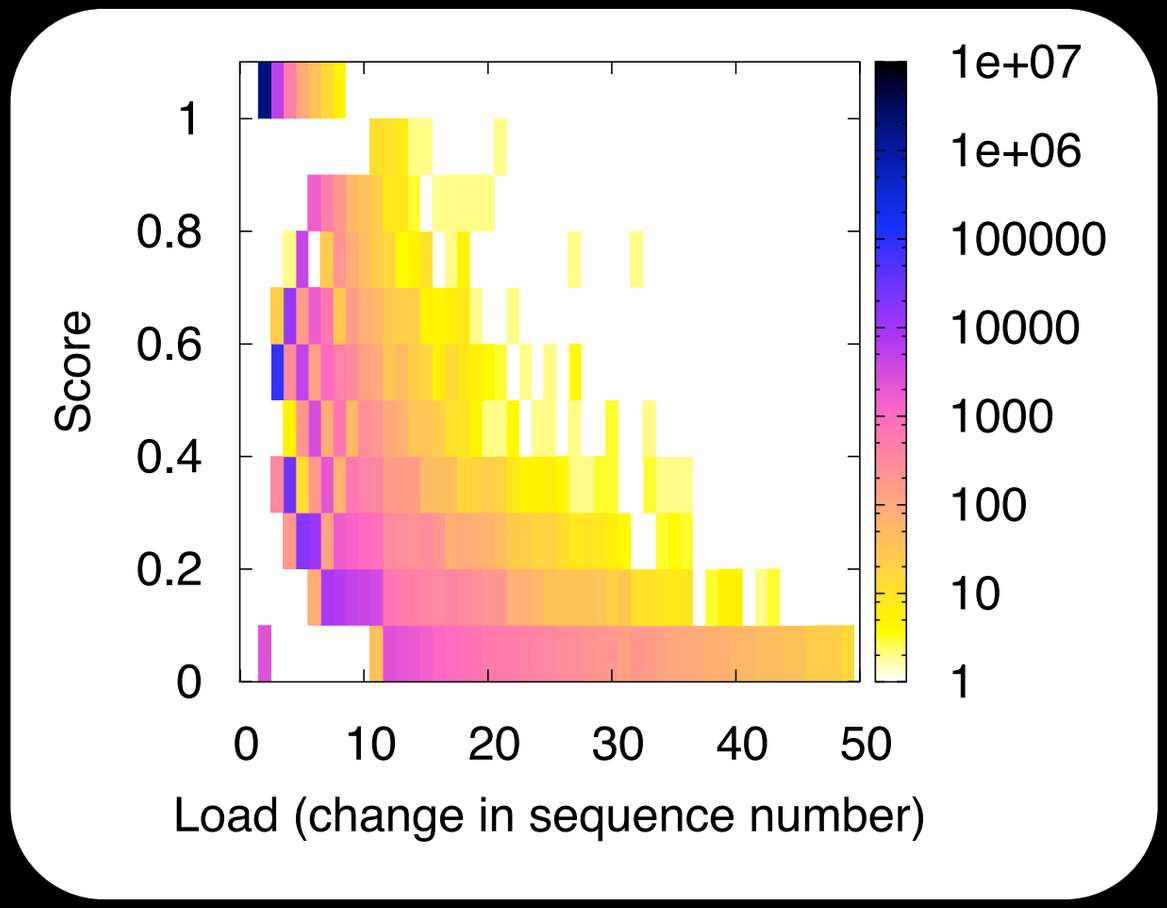

SIGCOMM 2004 AP

T-Fi plots focus on load

- Low load intervals are relegated to the left side

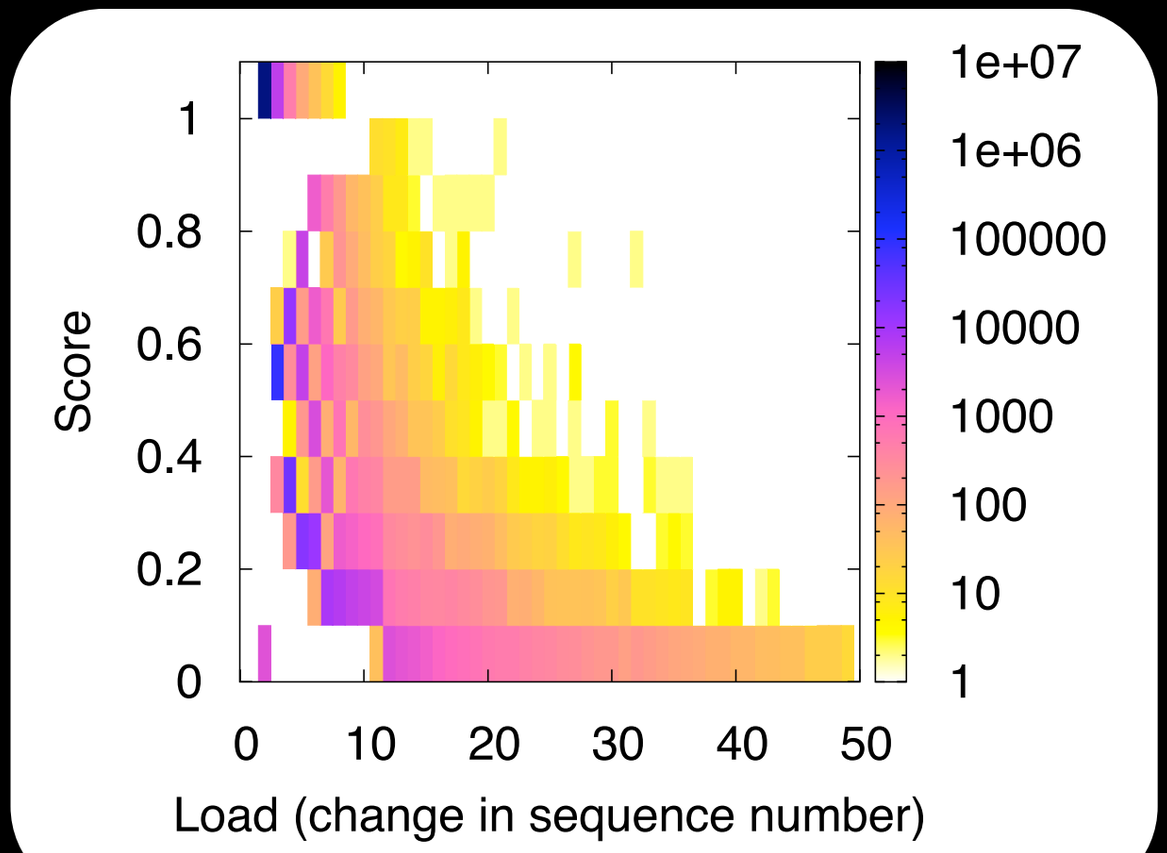

SIGCOMM 2004 AP

T-Fi plots focus on load

- Low load intervals are relegated to the left side

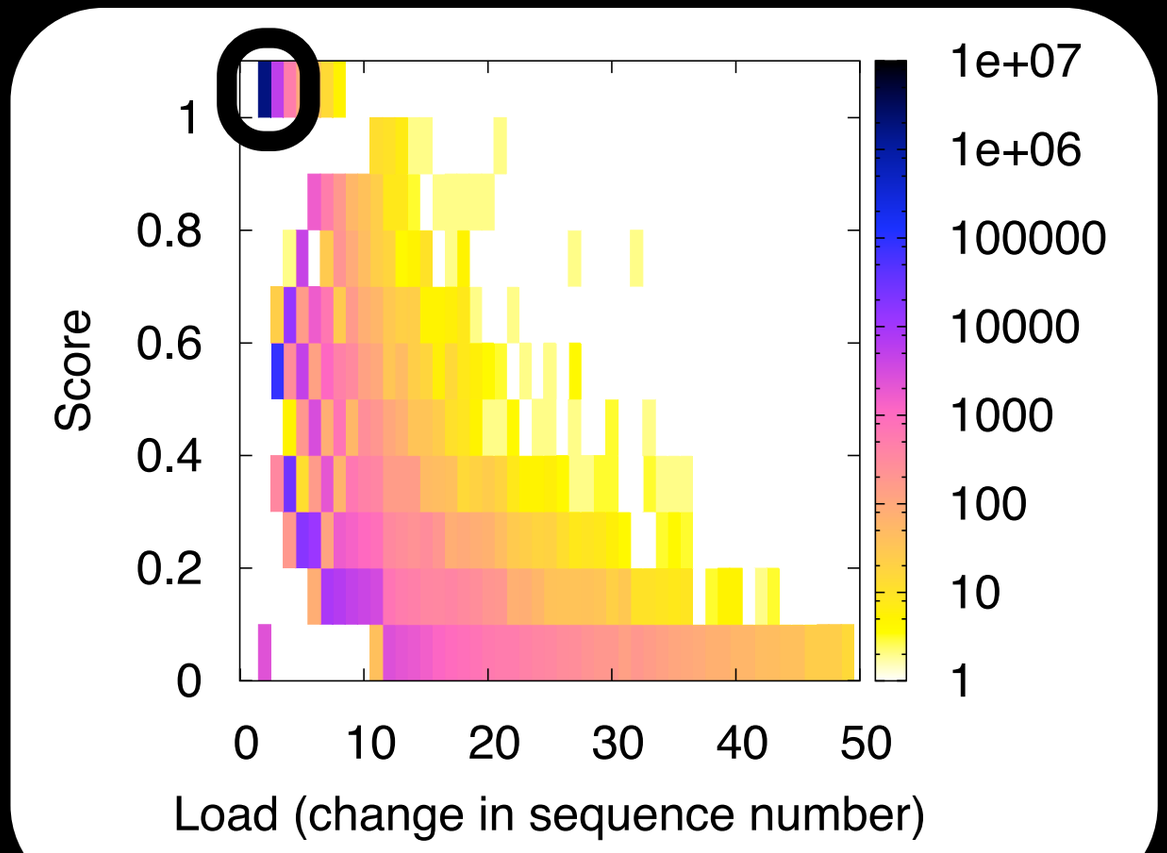

SIGCOMM 2004 AP

T-Fi plots focus on load

- Low load intervals are relegated to the left side

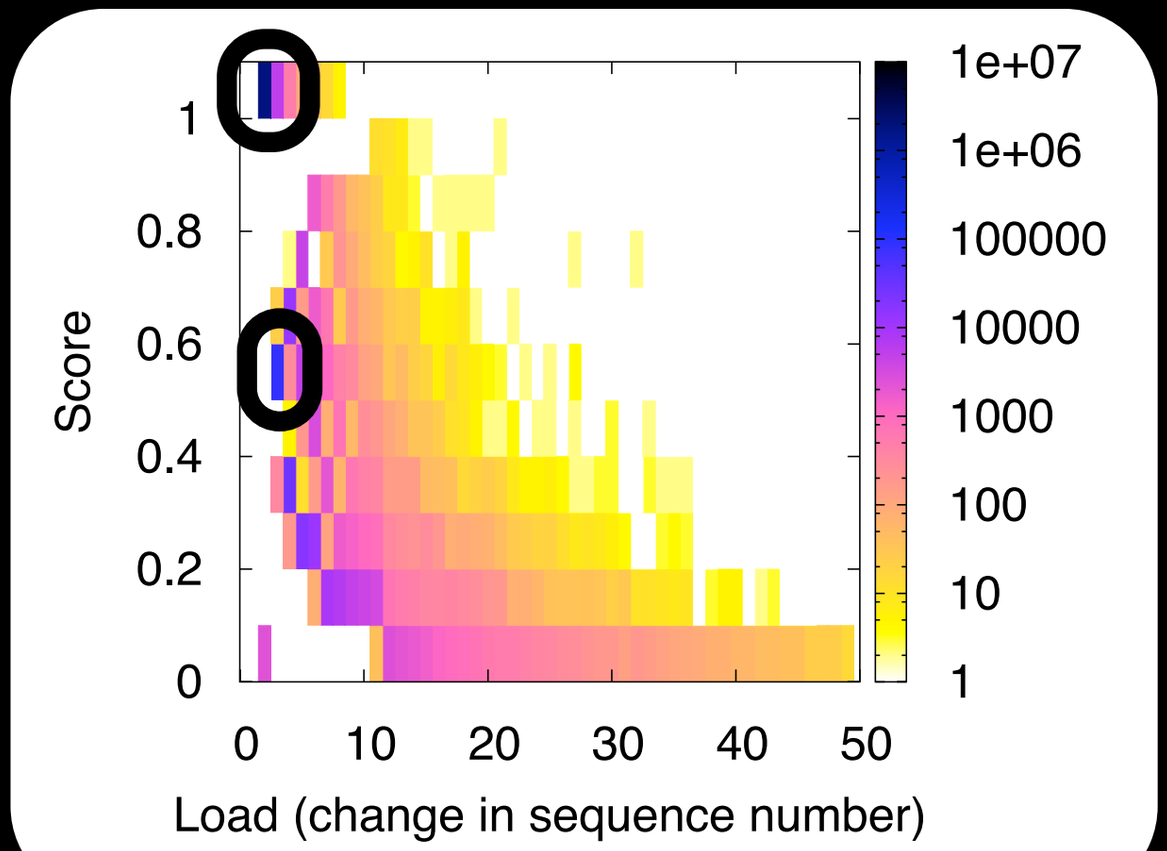

SIGCOMM 2004 AP

T-Fi plots focus on load

- Low load intervals are relegated to the left side

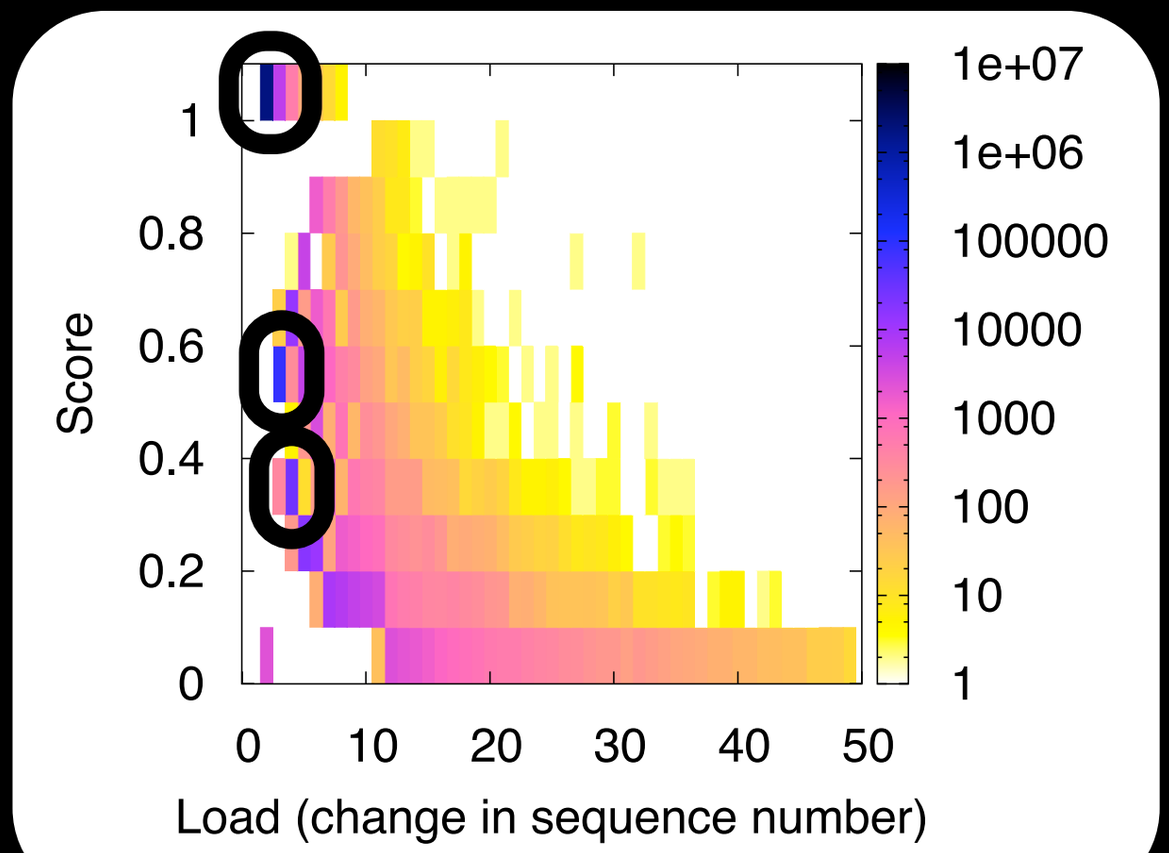

SIGCOMM 2004 AP

T-Fi plots focus on load

- Low load intervals are relegated to the left side

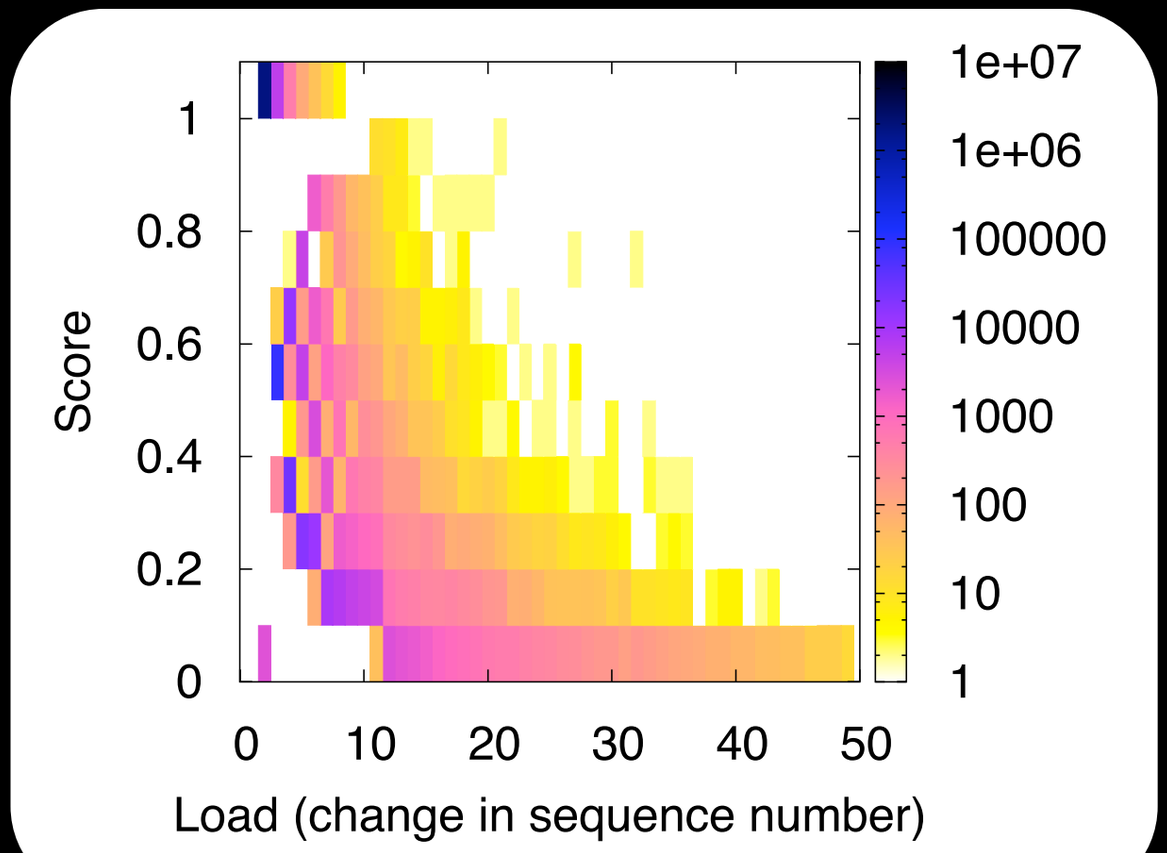

SIGCOMM 2004 AP

T-Fi plots focus on load

- Low load intervals are relegated to the left side
- High load intervals have low score

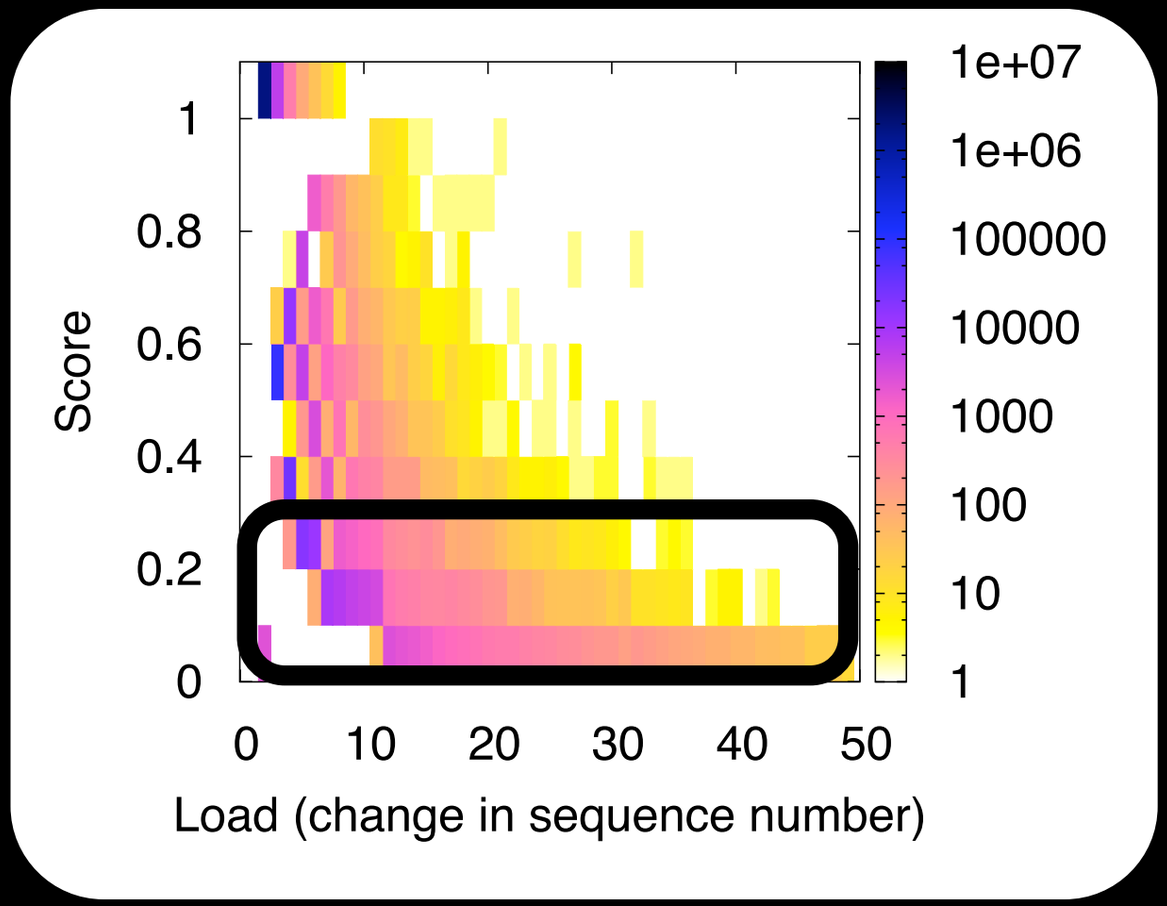

SIGCOMM 2004 AP

T-Fi plot comparison

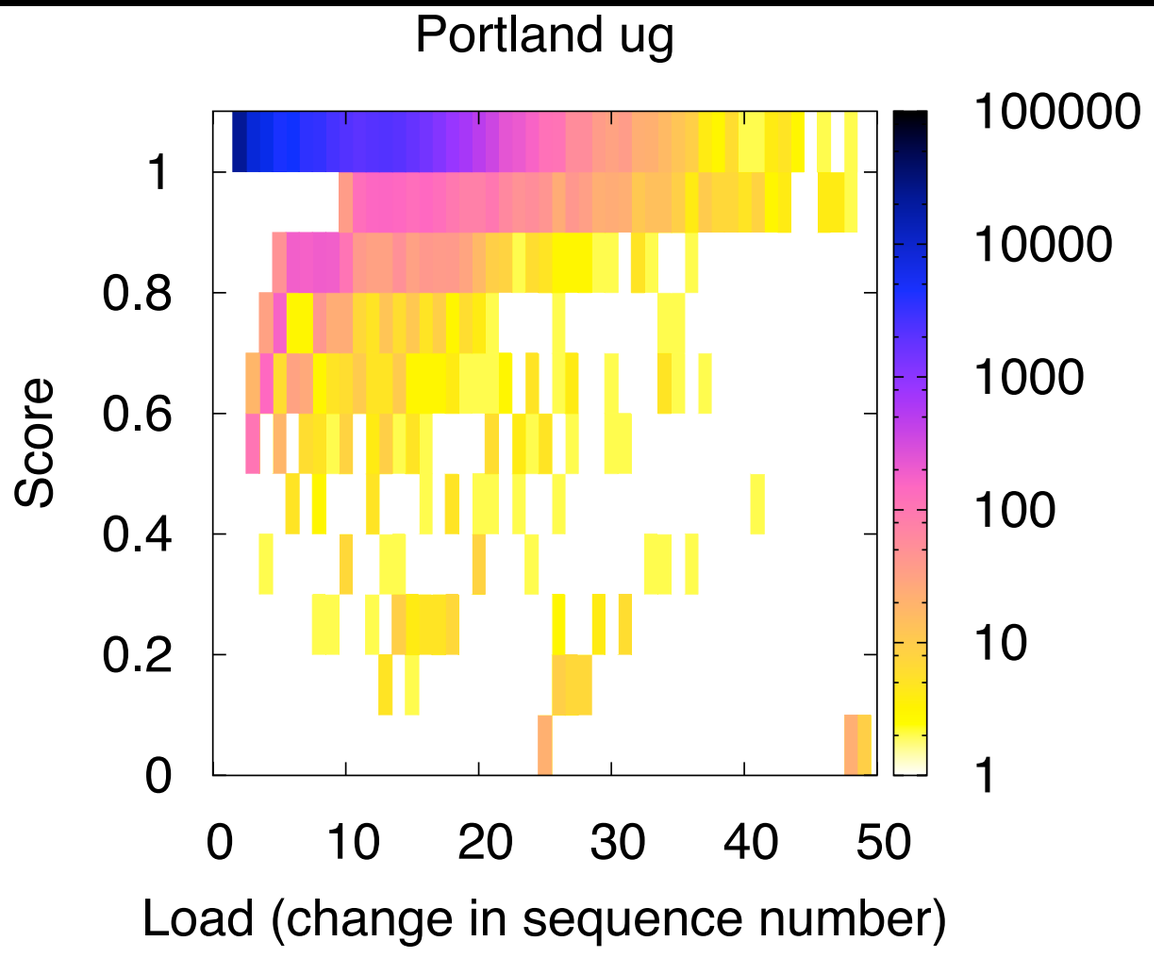

Portland PDX Dataset
Phillips et al

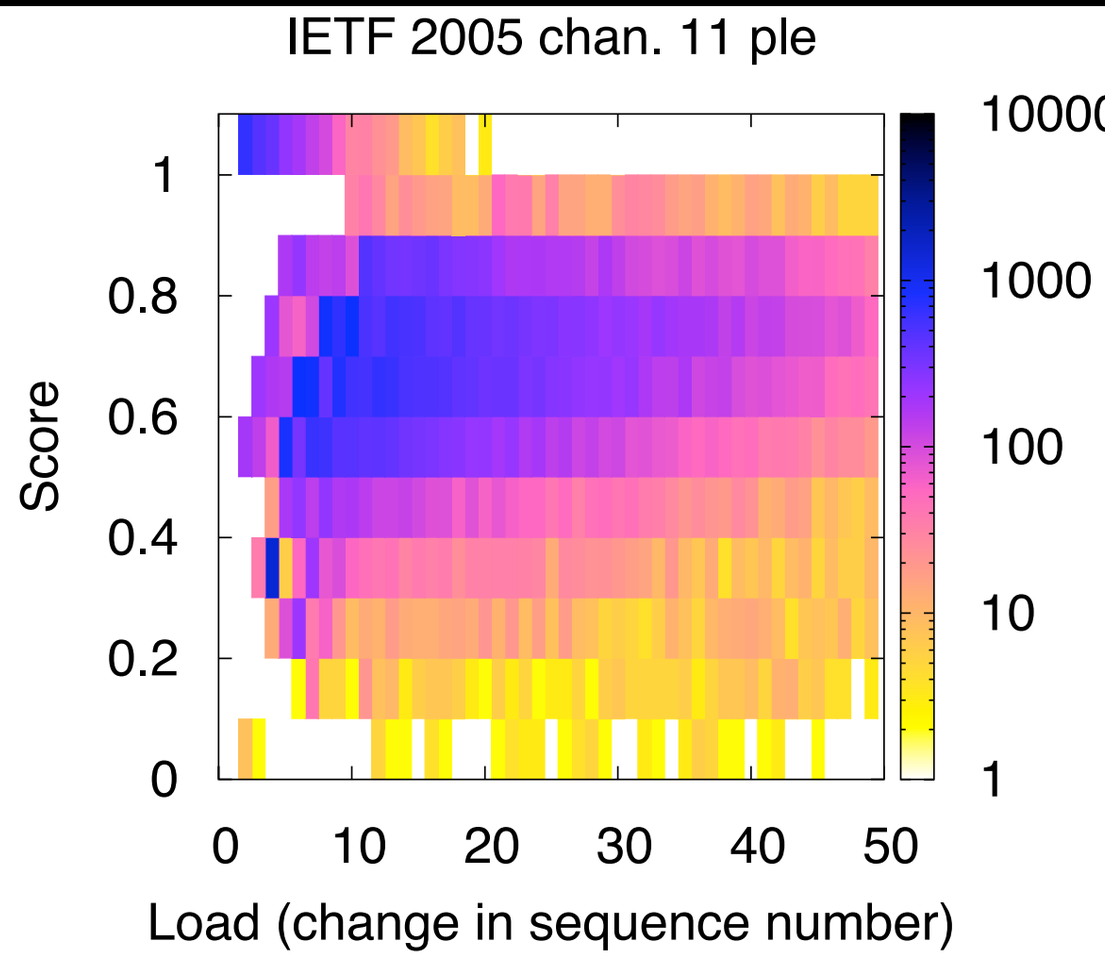

IETF 2005 Dataset
Jardosh et al

T-Fi plot comparison

1. Portland "ug" is more complete in 1 - 25 load intervals

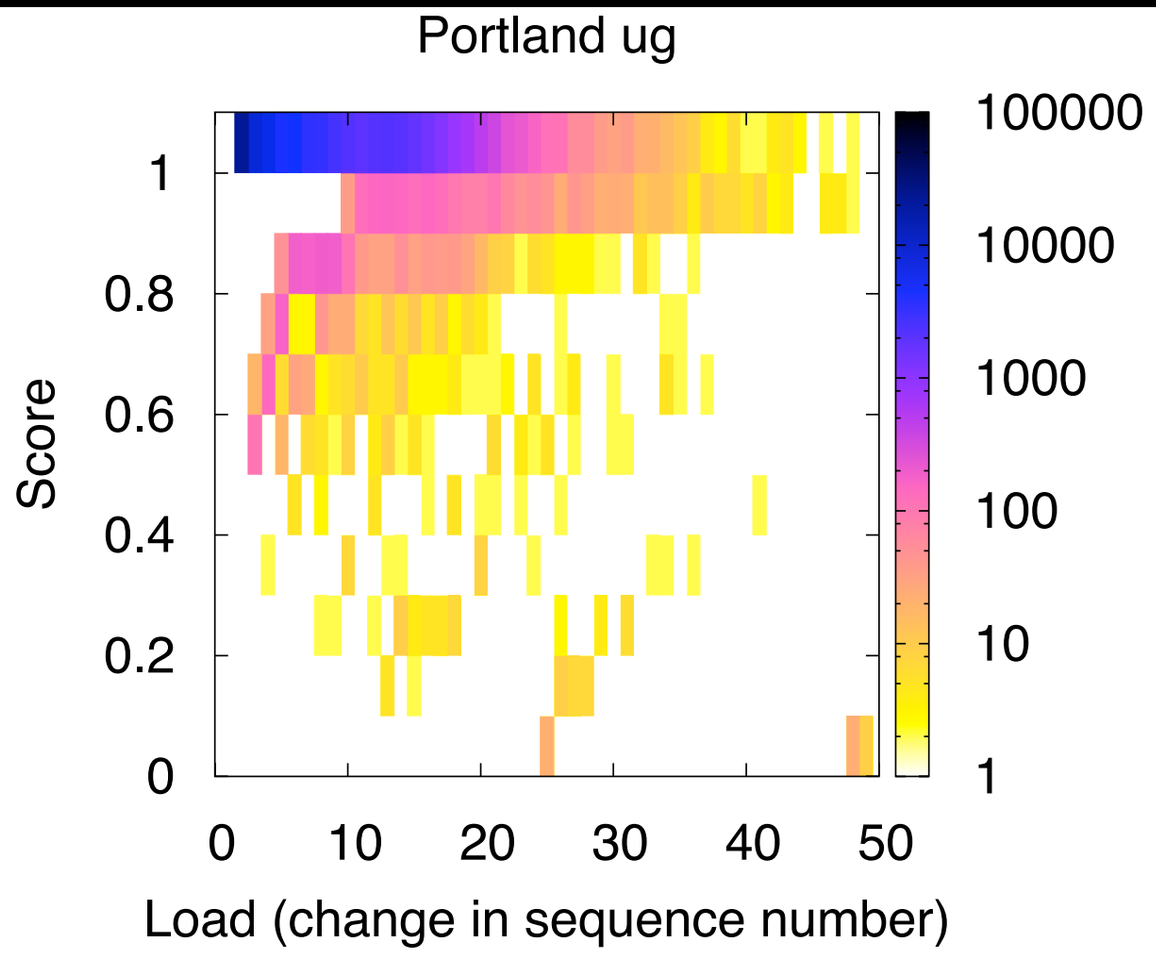

Portland PDX Dataset
Phillips et al

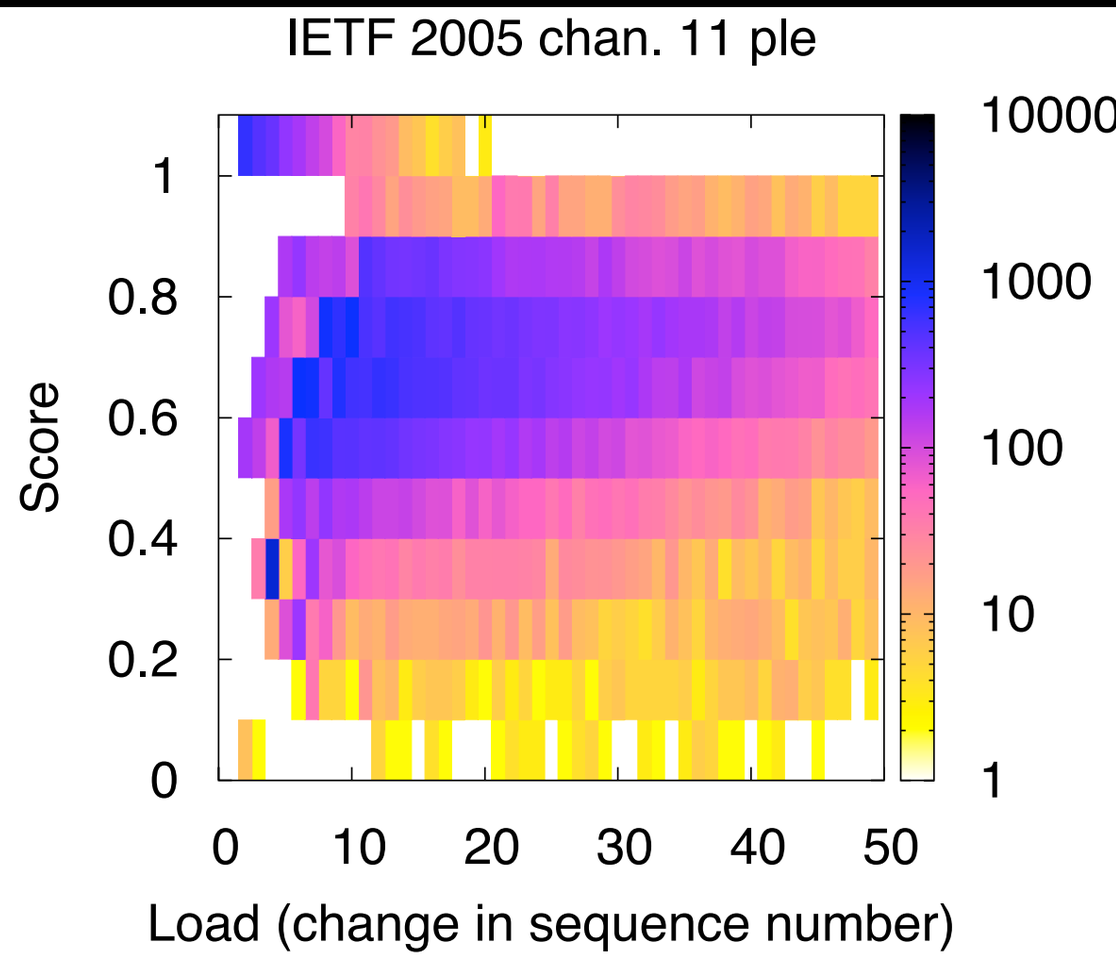

IETF 2005 Dataset
Jardosh et al

T-Fi plot comparison

1. Portland "ug" is more complete in 1 - 25 load intervals

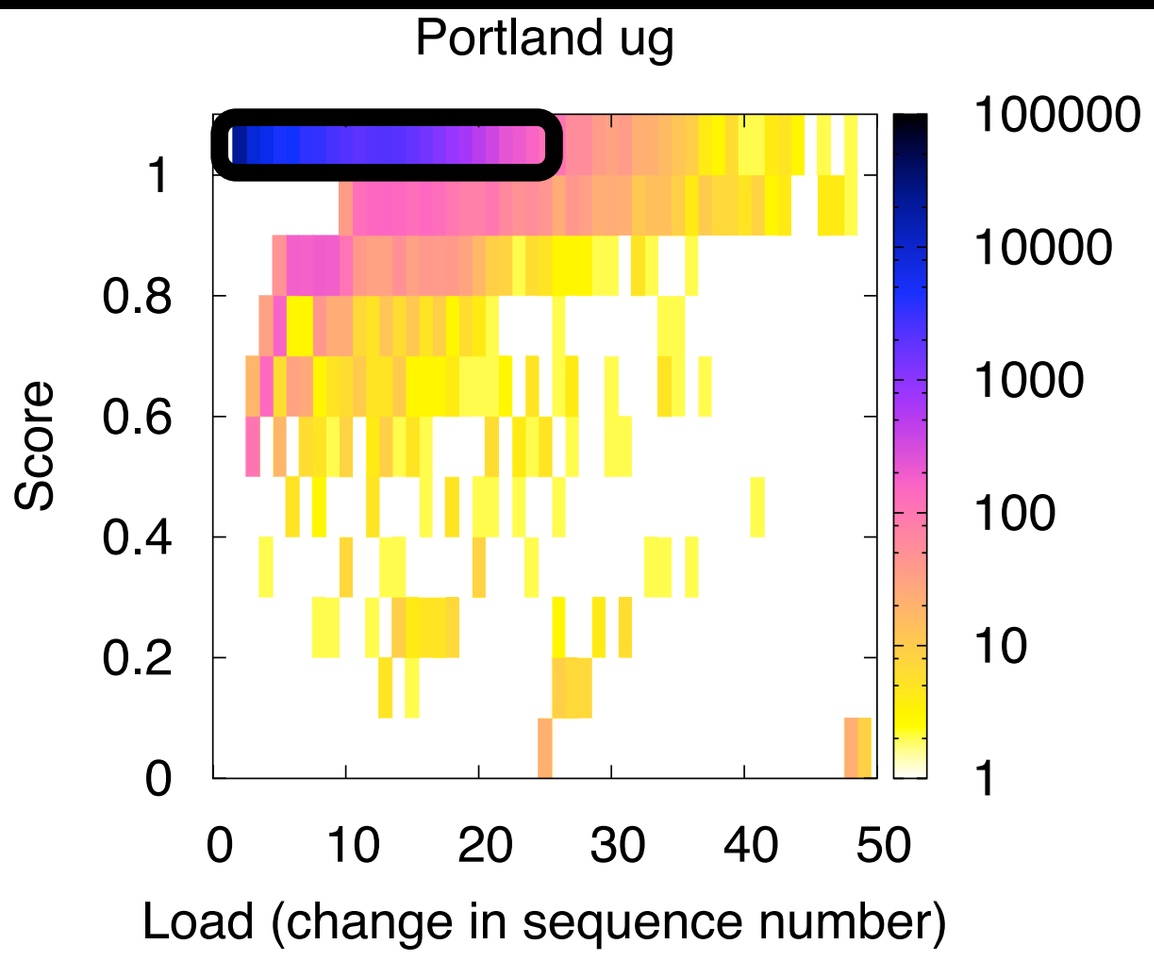

Portland PDX Dataset
Phillips et al

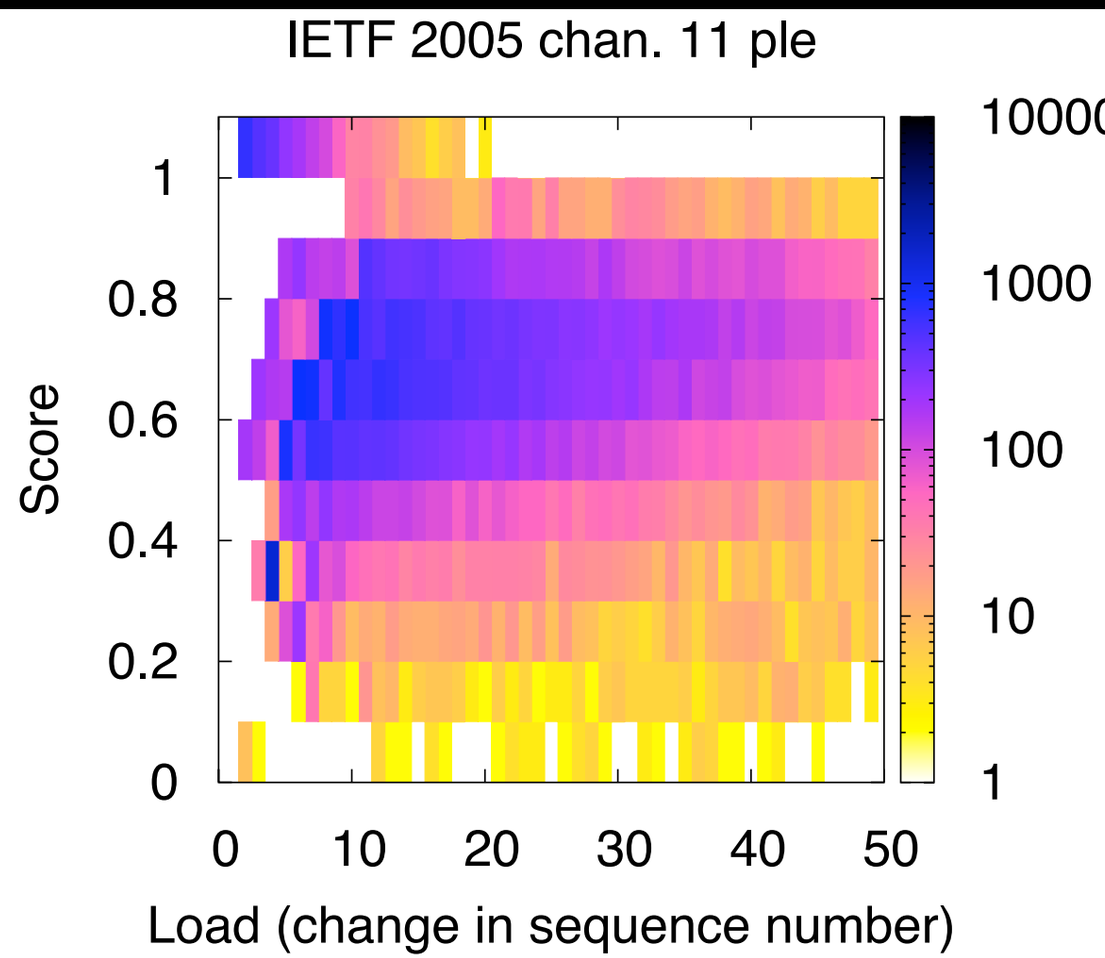

IETF 2005 Dataset
Jardosh et al

T-Fi plot comparison

1. Portland "ug" is more complete in 1 - 25 load intervals

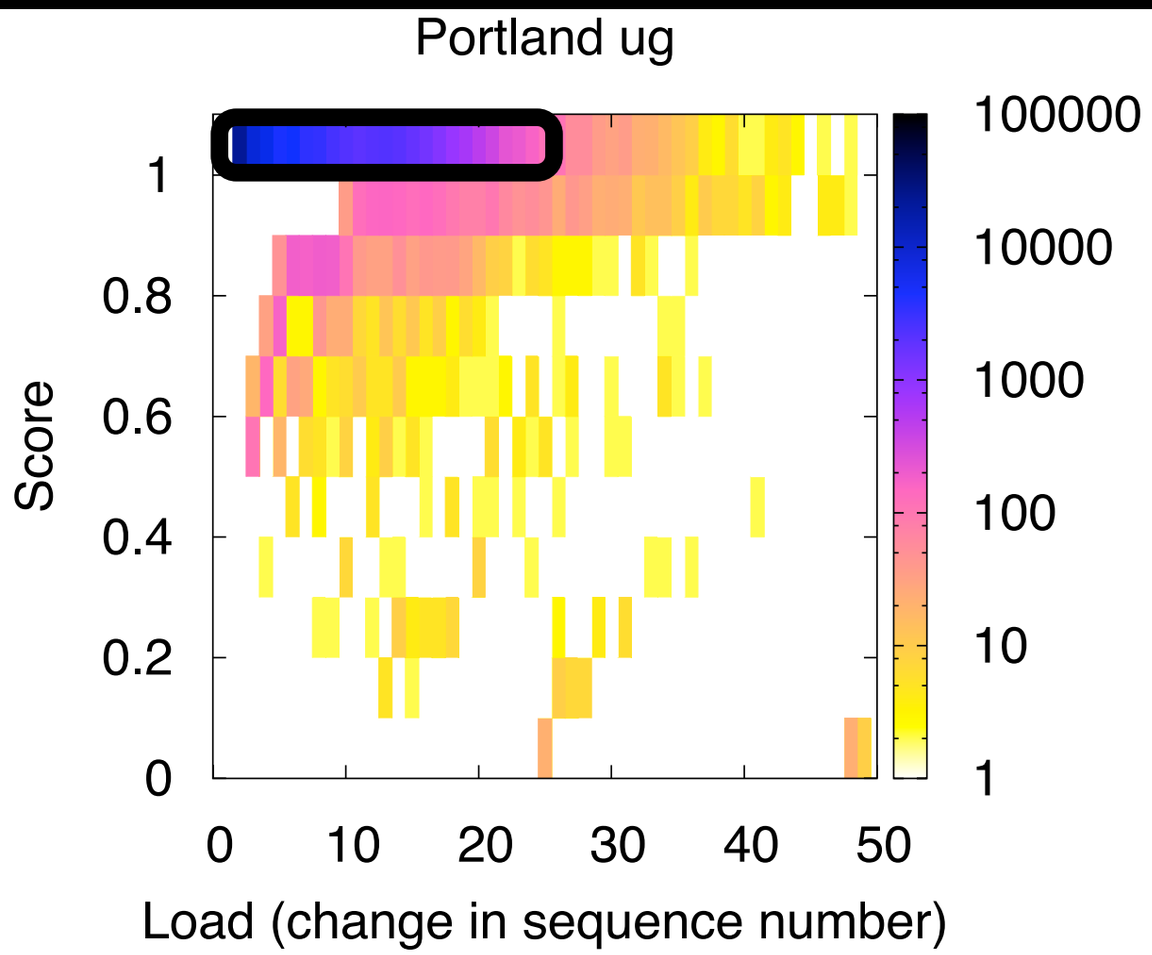

Portland PDX Dataset
Phillips et al

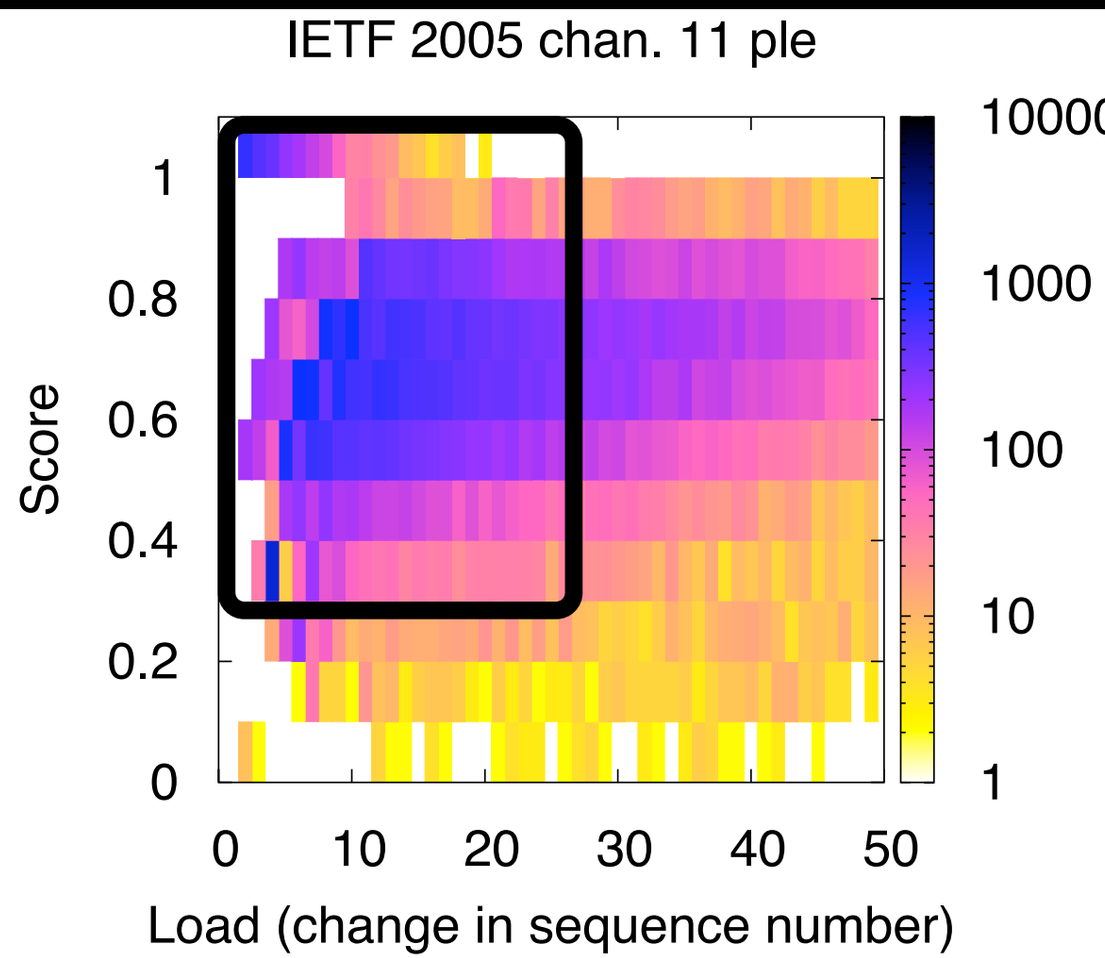

IETF 2005 Dataset
Jardosh et al

T-Fi plot comparison

1. Portland "ug" is more complete in 1 - 25 load intervals
2. IETF "chan. 11 ple" has more 30 - 50 load intervals

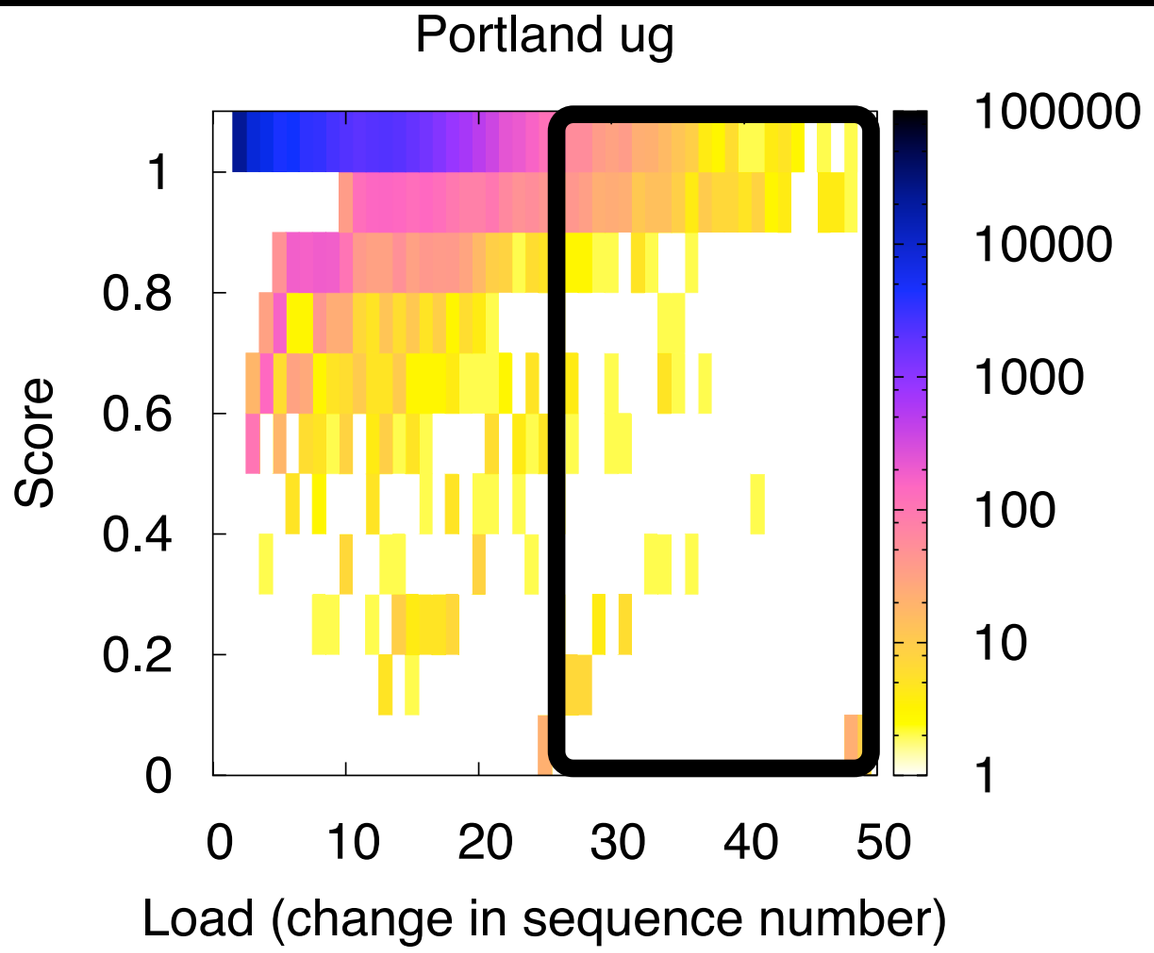

Portland PDX Dataset
Phillips et al

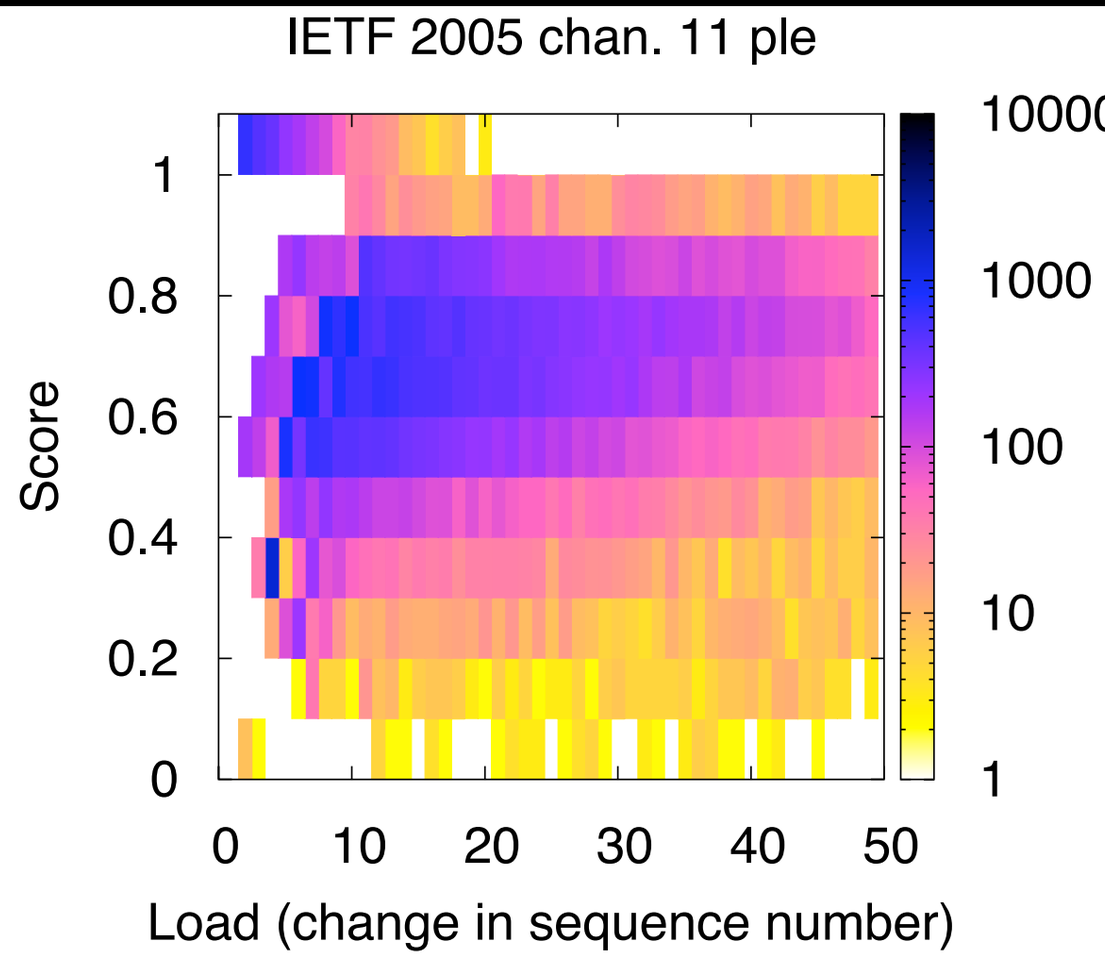

IETF 2005 Dataset
Jardosh et al

T-Fi plot comparison

1. Portland "ug" is more complete in 1 - 25 load intervals
2. IETF "chan. 11 ple" has more 30 - 50 load intervals

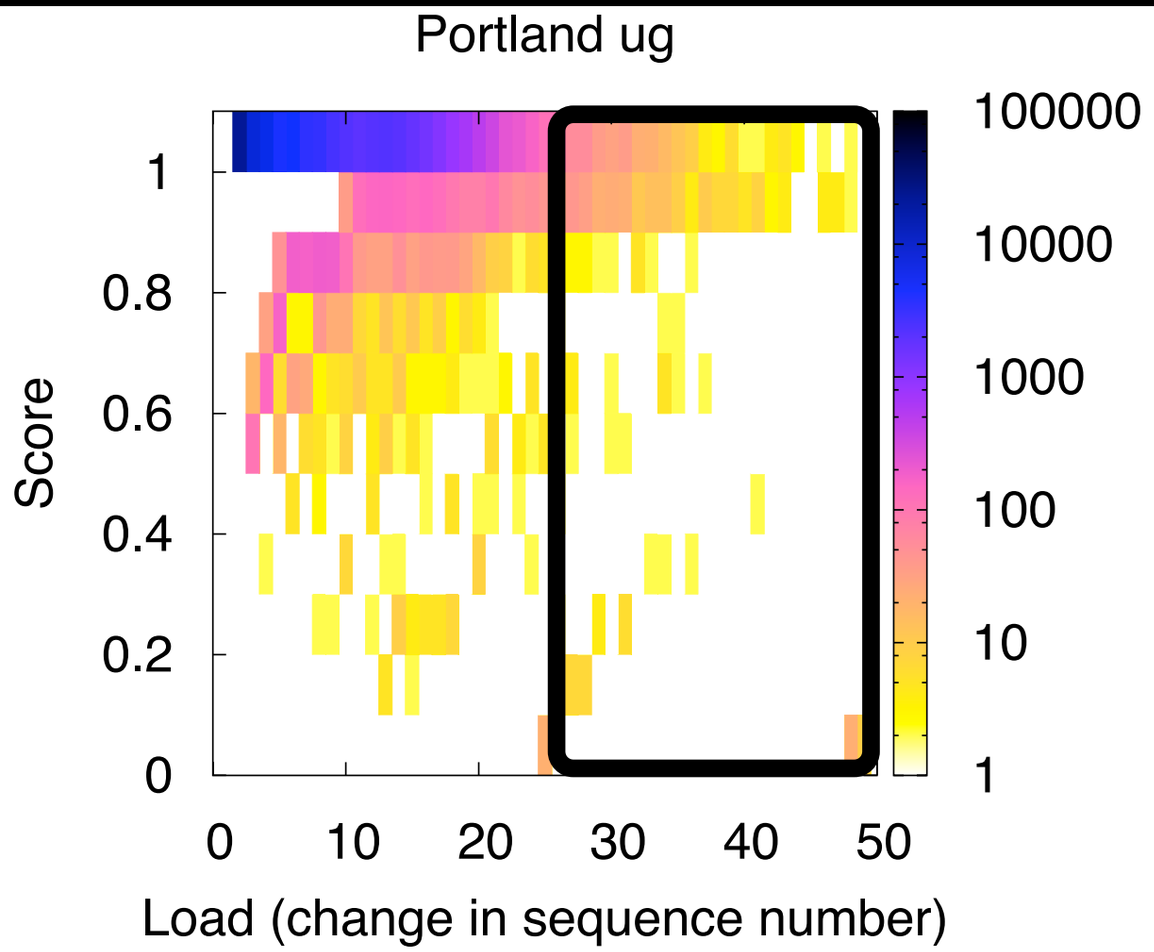

Portland PDX Dataset
Phillips et al

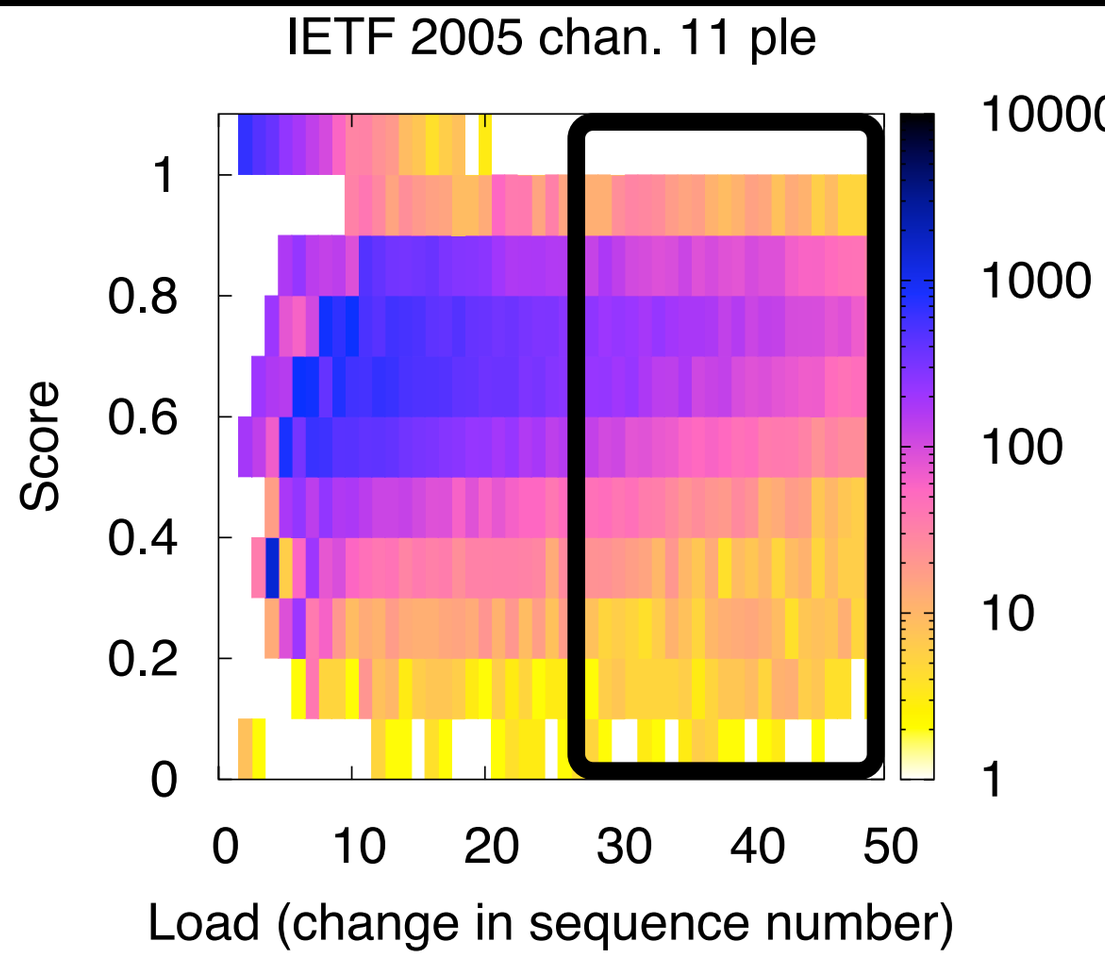

IETF 2005 Dataset
Jardosh et al

T-Fi plots

- T-Fi Plots can show other completeness measures
  - Completeness of a trace when there are many unique senders
    - Replace Load with # of unique senders

Trace Fidelity

Completeness

Did we capture all of the packets?

Accuracy

Did we timestamp the packets correctly?

Trace Fidelity

Completeness

Did we capture all of the packets?

T-Fi plots show trace completeness

Accuracy

Did we timestamp the packets correctly?

Accuracy

Did we timestamp the packets correctly?

Merging traces using packet timestamps

- Monitor applies timestamps to packets when it receives them
- Problem: Multiple monitors may not have synchronized clocks
- AP timestamps beacon packets before it sends them
- Solution: Synchronize monitors using beacon timestamps (Mahajan et al)

Synchronizing traces with beacon timestamps

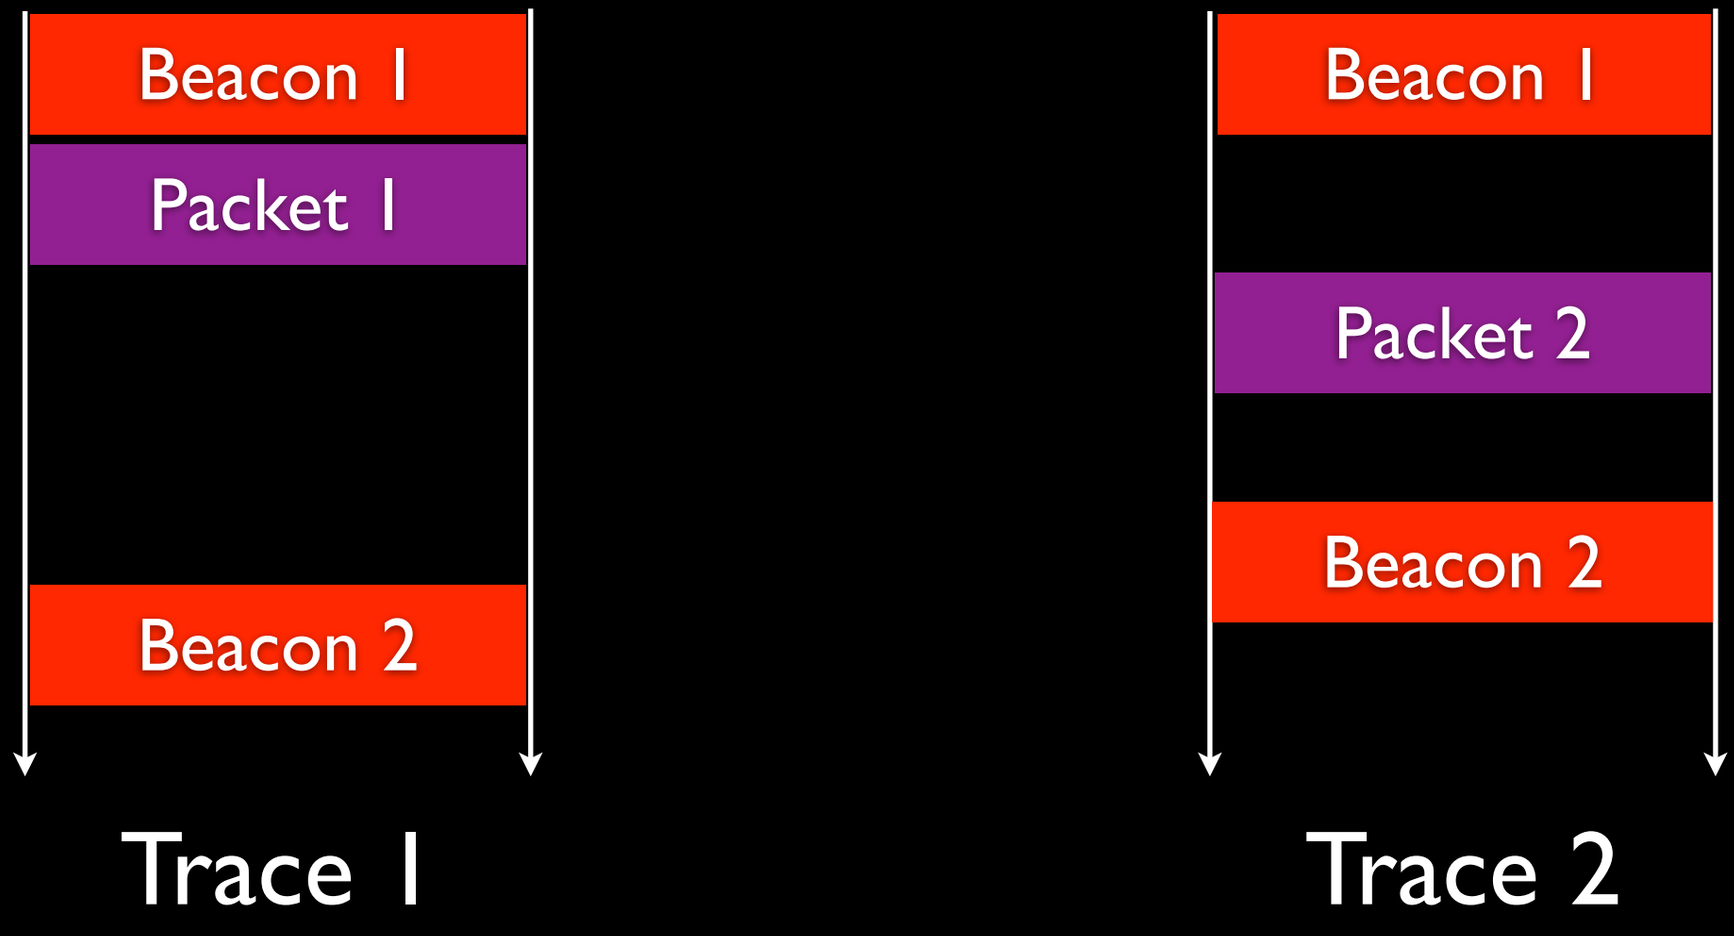

Synchronizing traces with beacon timestamps

Scale monitor timestamps to equal the interval from beacon timestamps

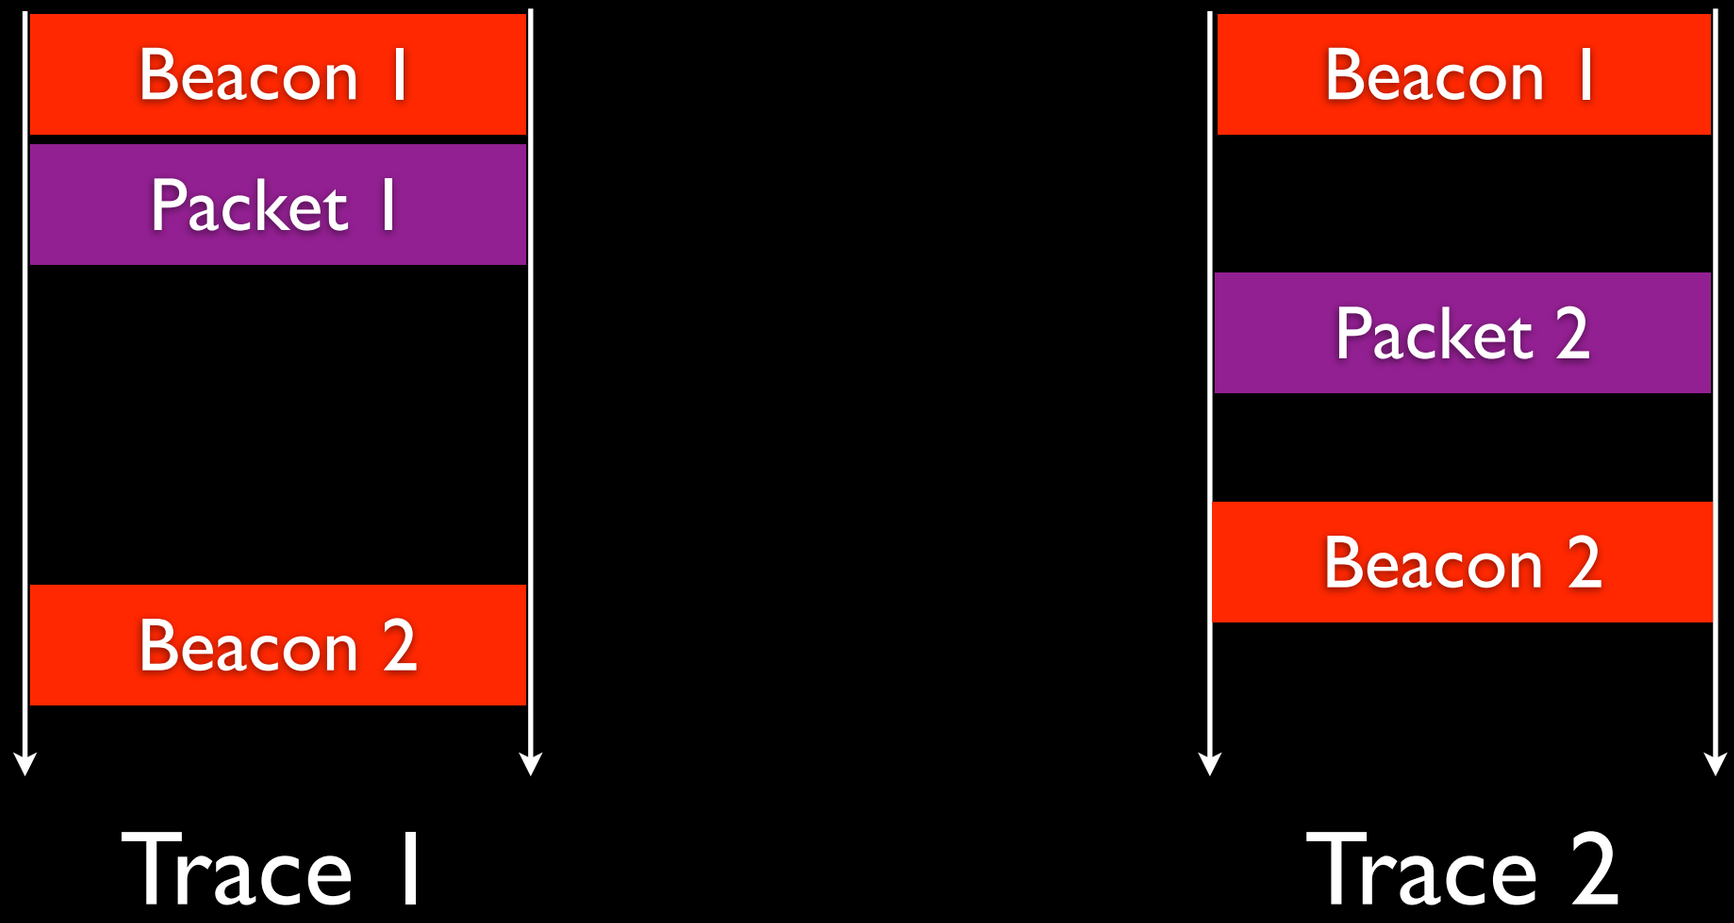

Synchronizing traces with beacon timestamps

Scale monitor timestamps to equal the interval from beacon timestamps

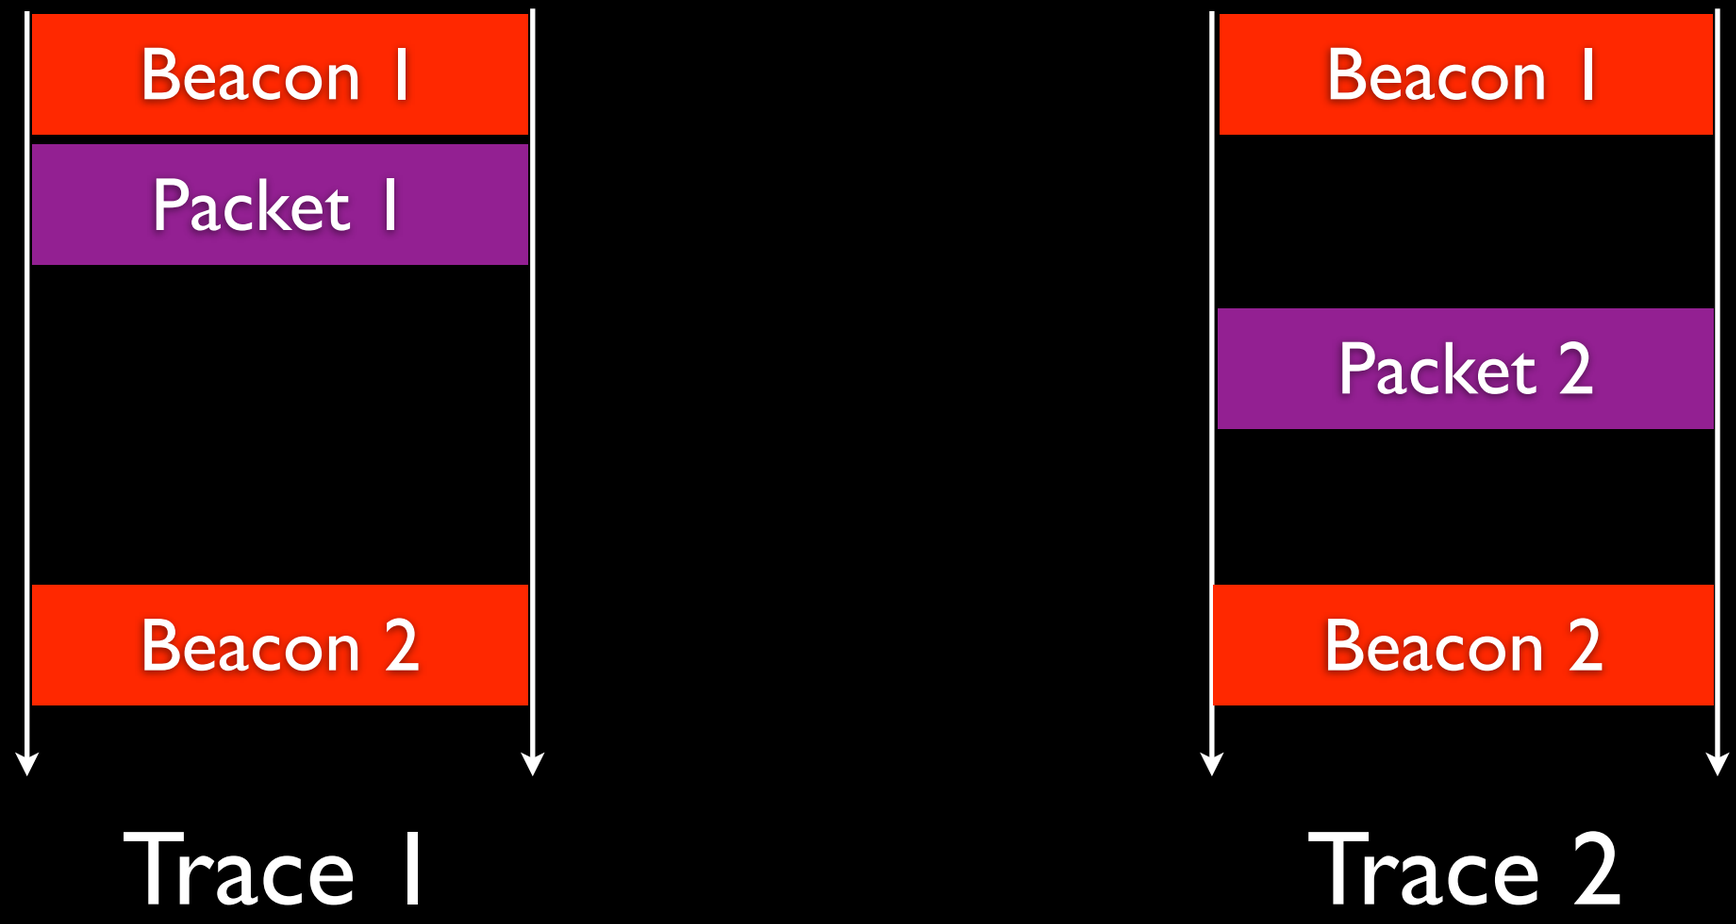

Synchronizing traces with beacon timestamps

Scale monitor timestamps to equal the interval from beacon timestamps

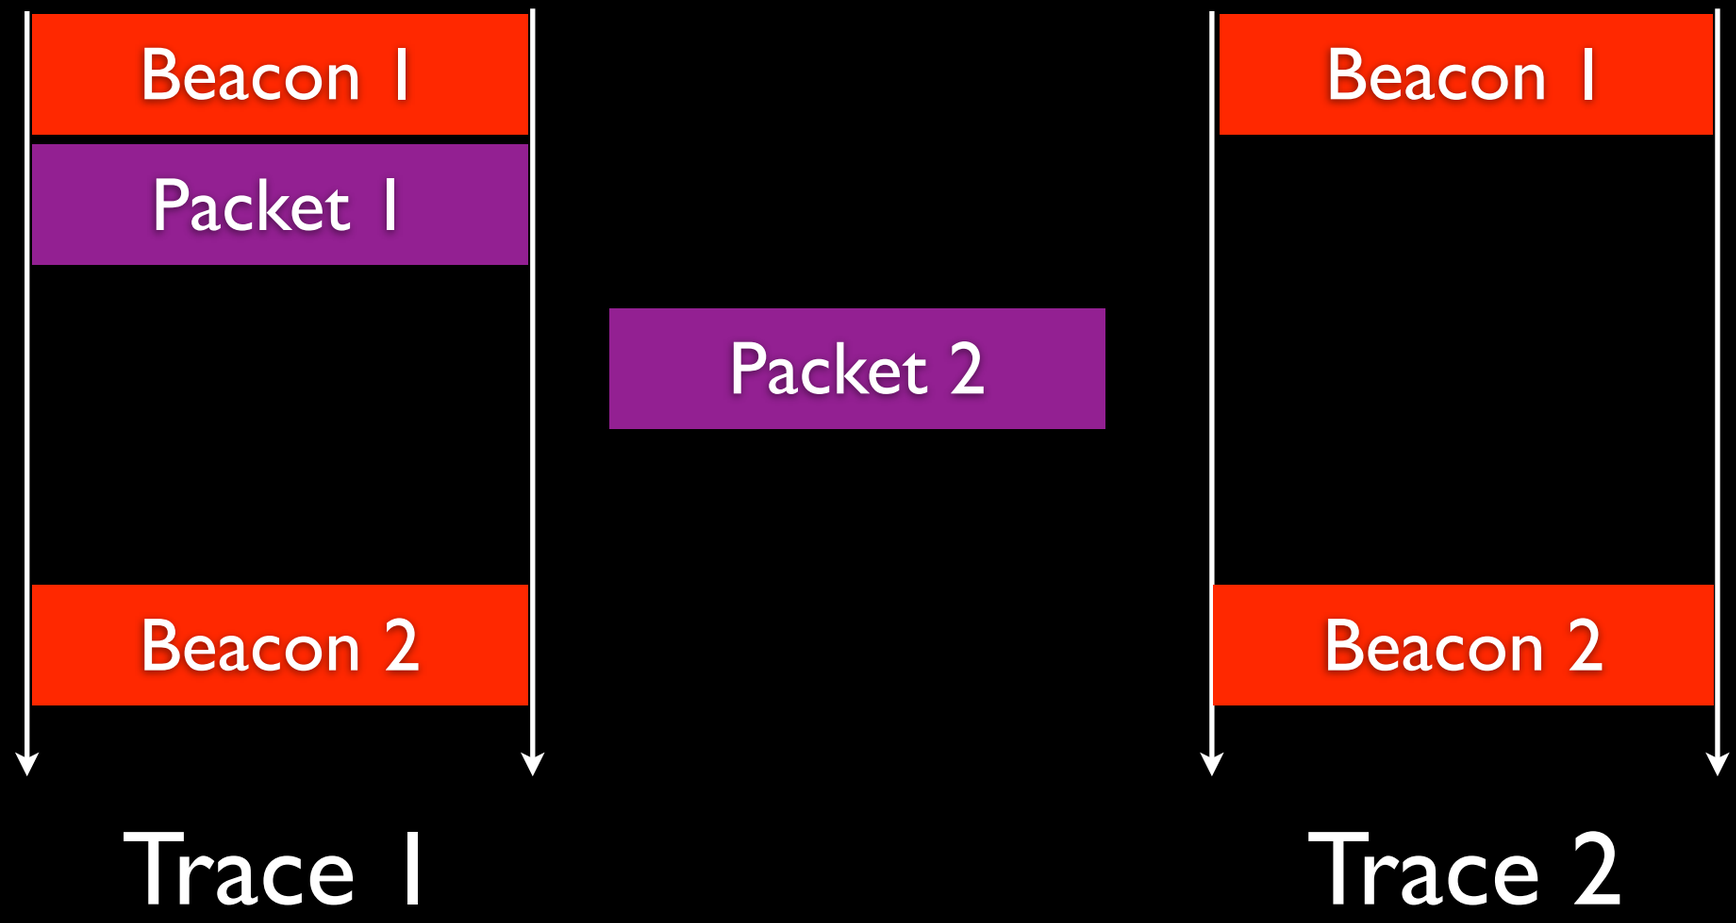

Synchronizing traces with beacon timestamps

Scale monitor timestamps to equal the interval from beacon timestamps

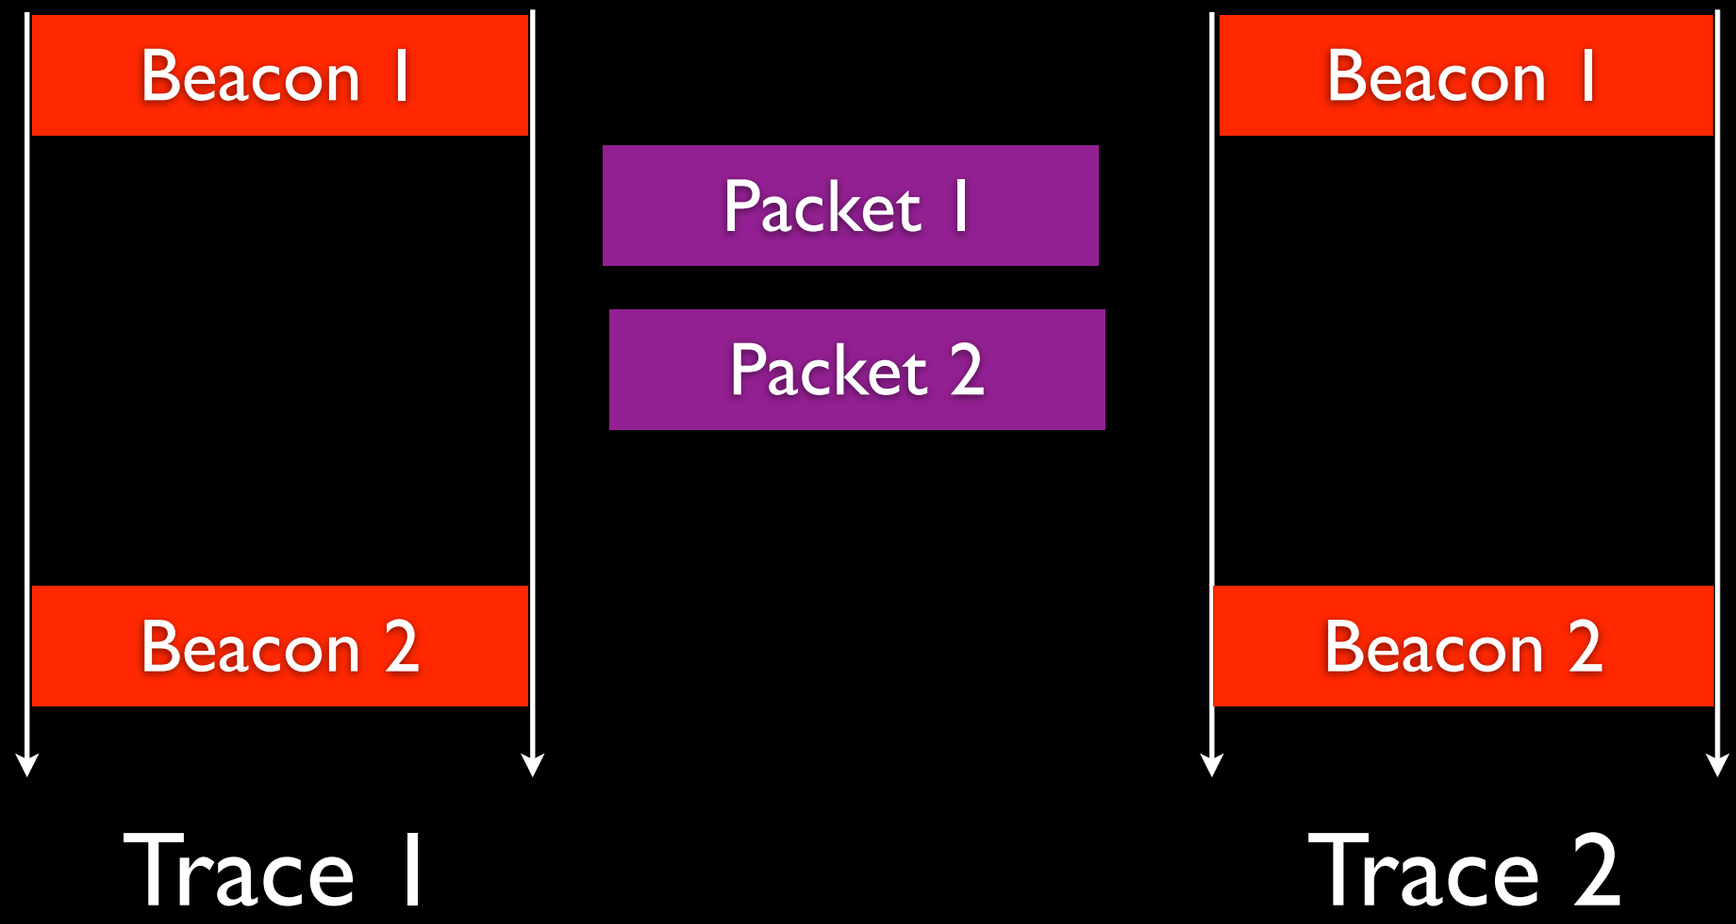

Compare monitor and beacon timestamps

- We measure the difference between beacon and monitor timestamps
- Is there clock skew at the monitor and/or AP?
- Clock diff. = Beacon Interval - Beacon Interval

Monitor AP

Accuracy is load-dependent

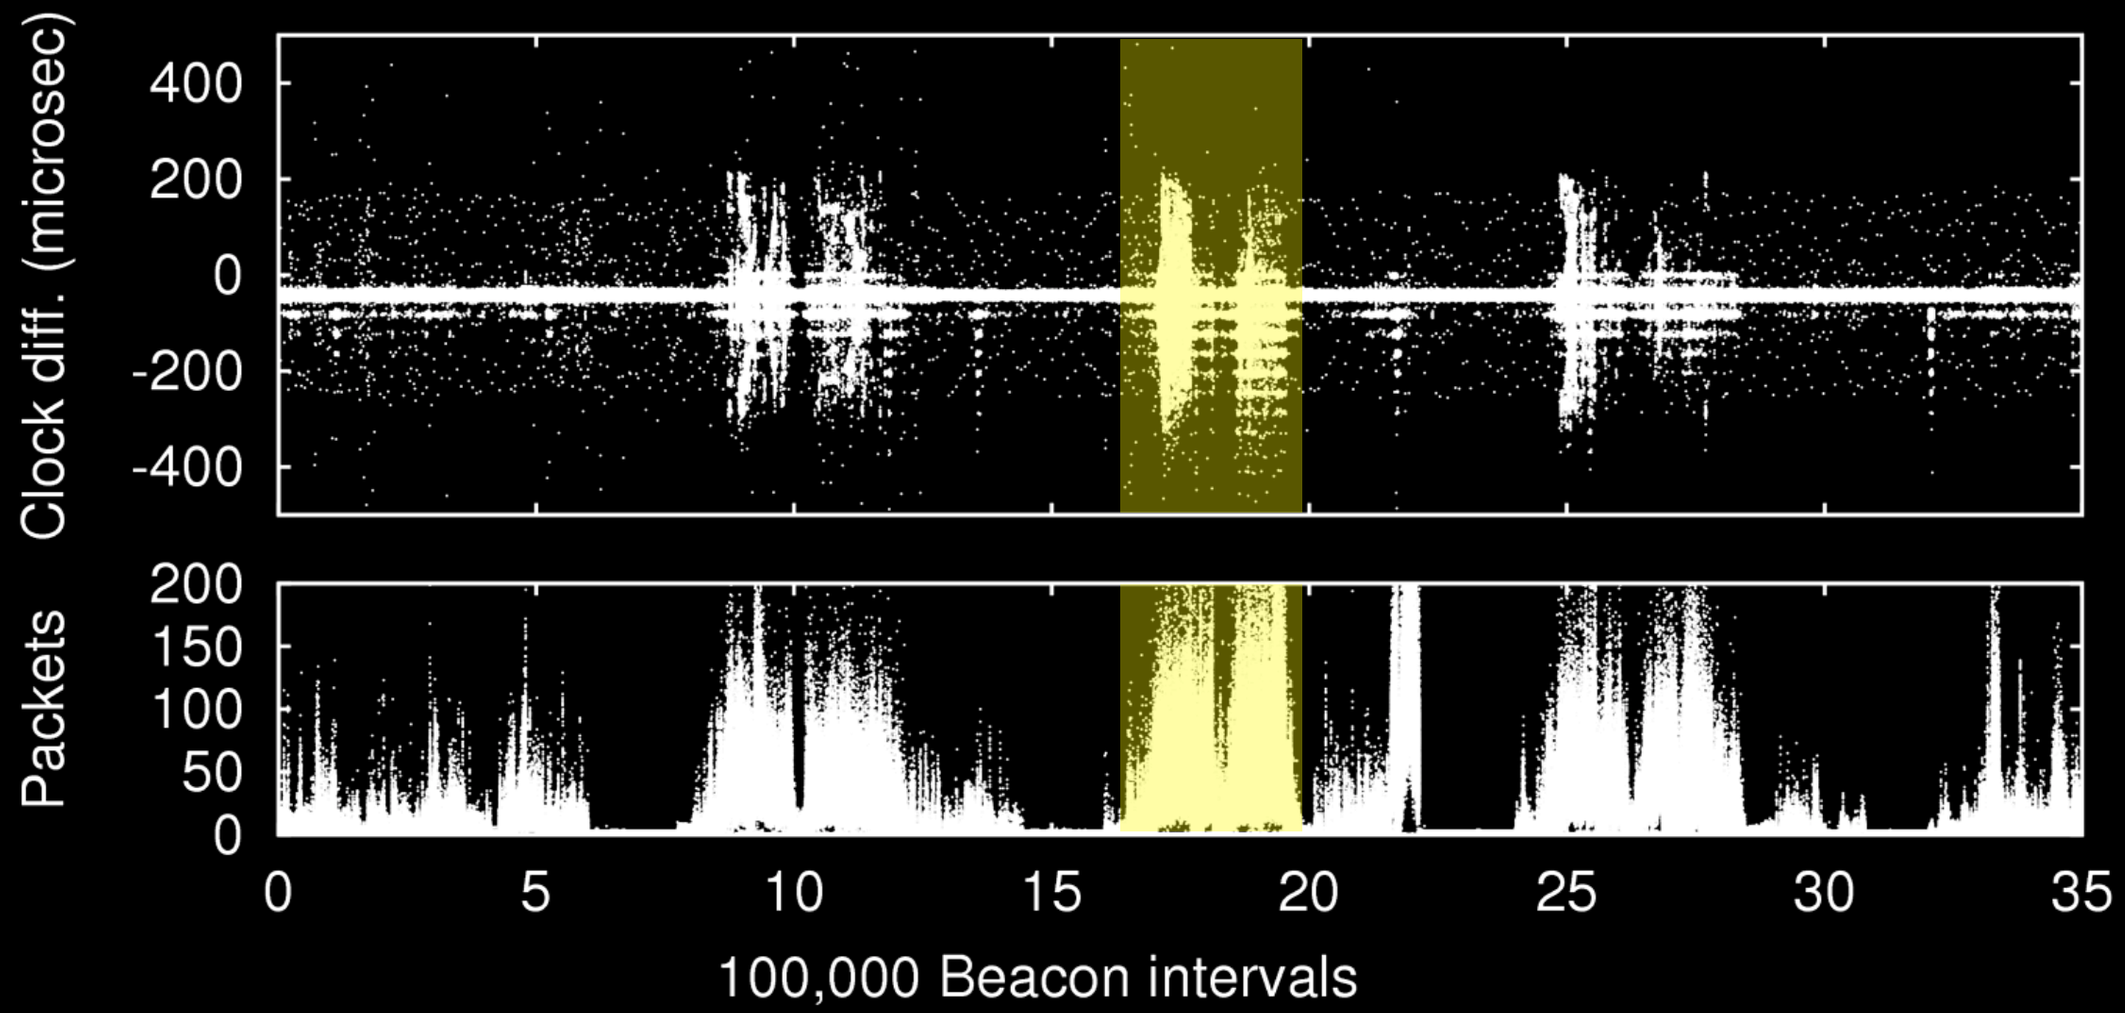

SIGCOMM 2004 Dataset
Rodrig et al.

Does clock difference exist inside beacon intervals?

Does clock difference exist inside beacon intervals?

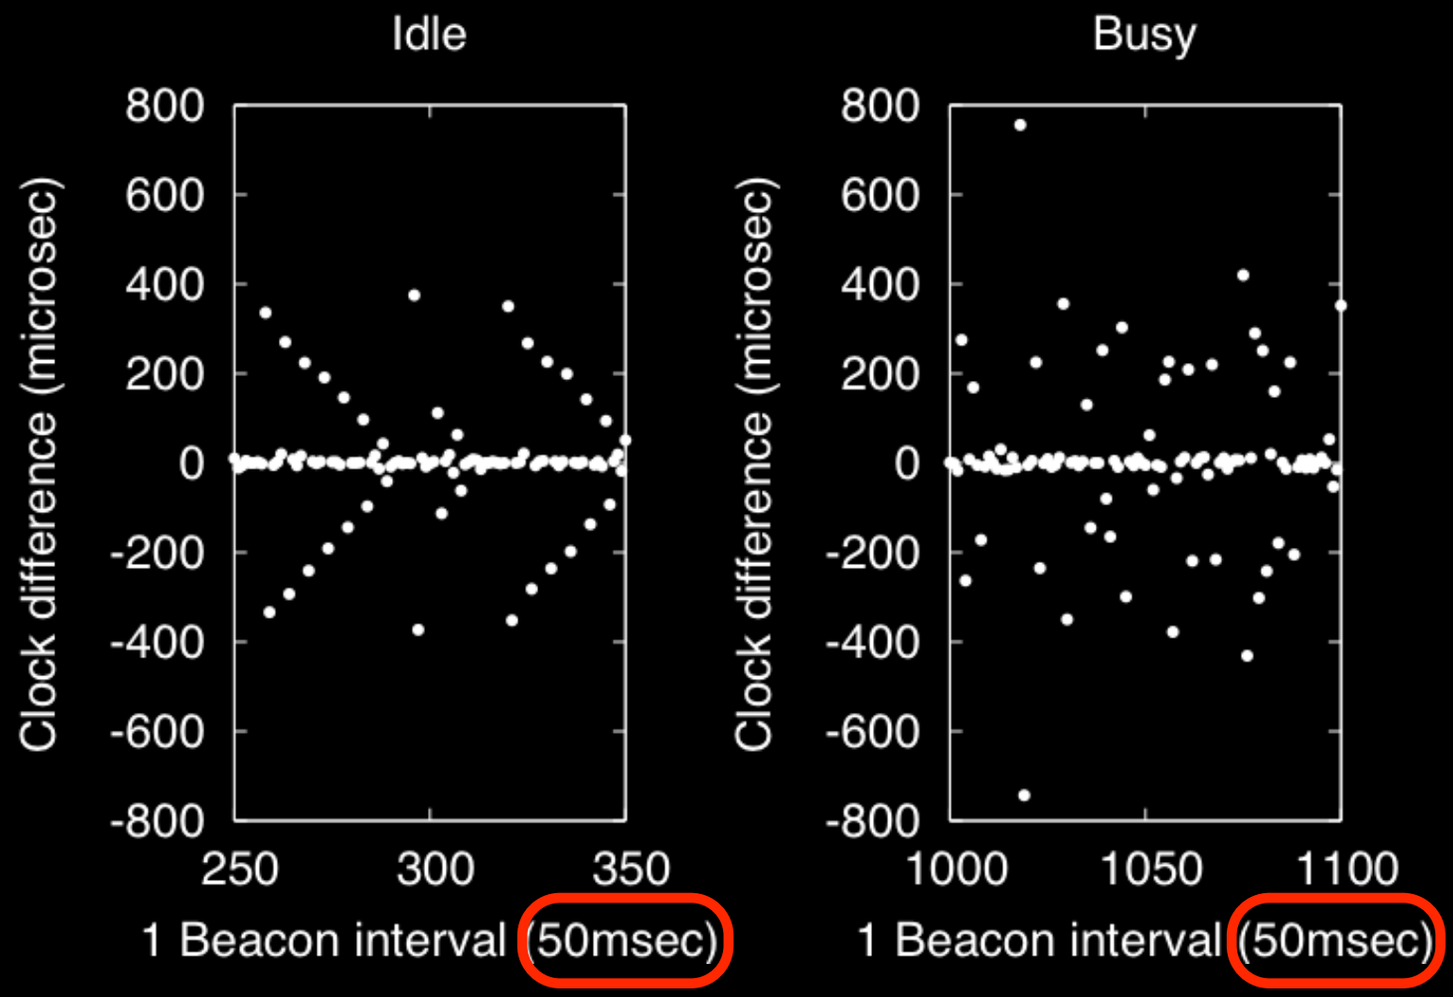

Does clock difference exist inside beacon intervals?

Significant clock differences can exist inside 100ms intervals

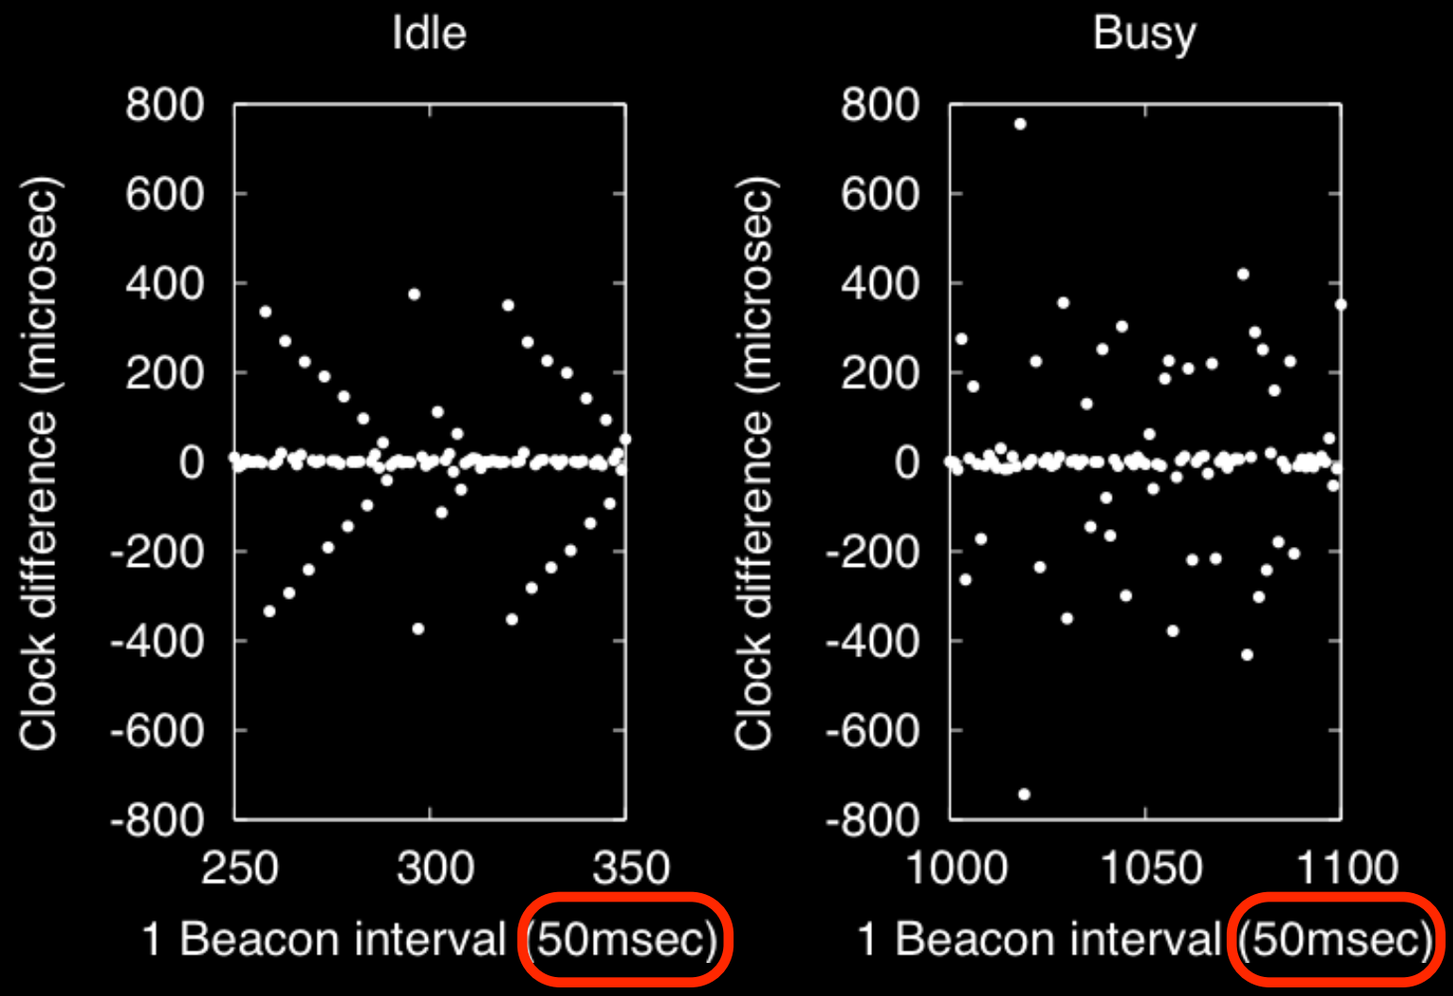

Trace Fidelity

Completeness

Did we capture all of the packets?

Accuracy

Did we timestamp the packets correctly?

Trace Fidelity

Completeness

Did we capture all of the packets?

T-Fi plots show trace completeness

Accuracy

Did we timestamp the packets correctly?

Trace Fidelity

Completeness

Did we capture all of the packets?

T-Fi plots show trace completeness

Accuracy

Did we timestamp the packets correctly?

Load increases frequency of timestamp error

Trace Fidelity

Completeness

Did we capture all of the packets?

T-Fi plots show trace completeness

Accuracy

Did we timestamp the packets correctly?

Load increases frequency of timestamp error

Merging algorithms have a faulty assumption

Conclusions

- Completeness and accuracy depend on load
- The fundamental assumption behind merging algorithms is flawed
- Future Work: Identifying the fidelity of a trace in real-time

http://www.cs.umd.edu/projects/wifidelity