1-Bit Spinning

Measurement Results

Marcus Ihlar - Ericsson

Scope

- How good can we get with single bit measurements?

- What is the impact of different types of network impairments?

Measurement System

```
┌─────────────────┐        ┌─────────────────┐        ┌─────────────────┐
│                 │ ◄─10ms─►│    Passive      │ ◄─20ms─►│                 │
│   QUIC Server   │         │  measurement    │  10mbit │   QUIC Client   │
│                 │         │     device      │         │                 │
└─────────────────┘        └─────────────────┘        └─────────────────┘
```

Loss and reordering configurable

- RTT samples collected at measurement device and QUIC Server, full e2e RTT measured.

- Server reports every RTT sample it measures internally and samples at the frequency of spin transitions.

- Measurement device uses simple heuristics to reject bad samples:
  - Record number of packets between transitions
  - If the number of packets with new spin >= 10% of previous packet count, then the edge is validated
 - Reject samples obviously too small (1/8 of initial RTT).

3

Findings

- 1 bit with simple heuristics gives "good enough" measurments.

- Downstream impairments have minimal effect on spin bit mechanism.

- Reordering effects are mitigated with use of simple heuristics.

- Uniformly random packet loss has very small impact on spin bit mechanism.

  – Slight bias towards larger values in spin measurements.

- Lower sampling frequency of spin bit compared to server does not seem to have large impact on measurement results.

  – Slight bias towards lower values in high frequency measurements.

RTT Estimates - reordering

5% Reorder, depth 10ms
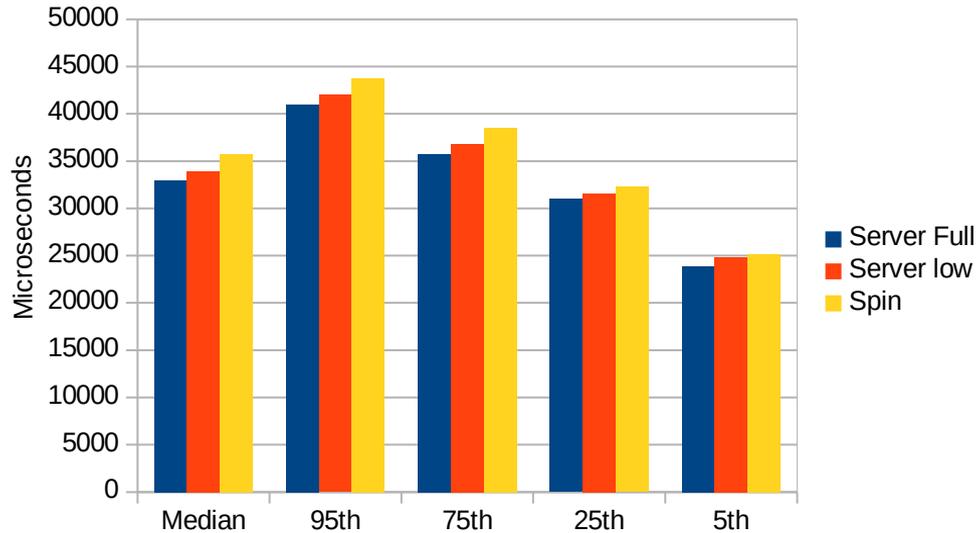
10% Reorder, depth 10ms
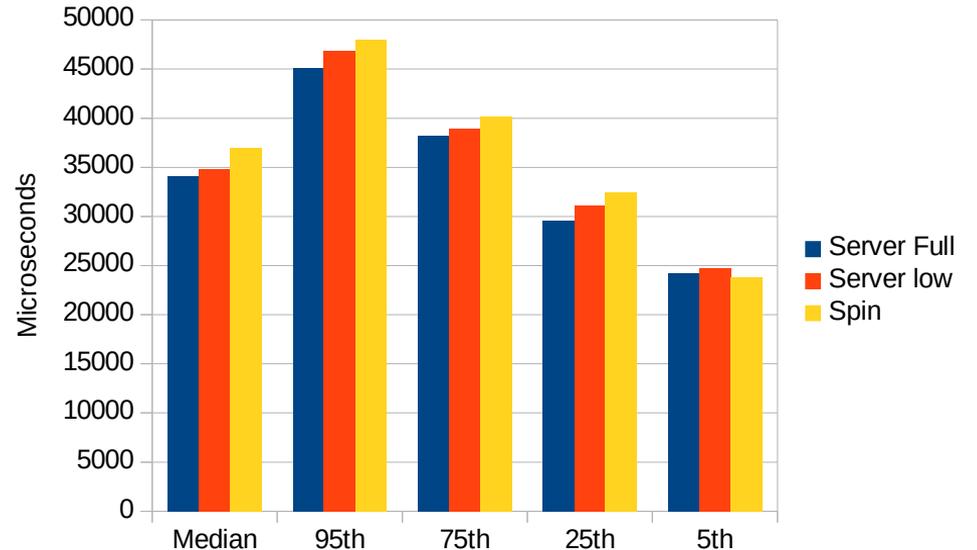
Number of RTT samples

5% Reorder, depth 10ms

Sample size	
Server high	2970 (100%)
Server low	792 (27%)
Spin	596 (20%)

196 samples rejected due to heuristics

10% Reorder, depth 10ms

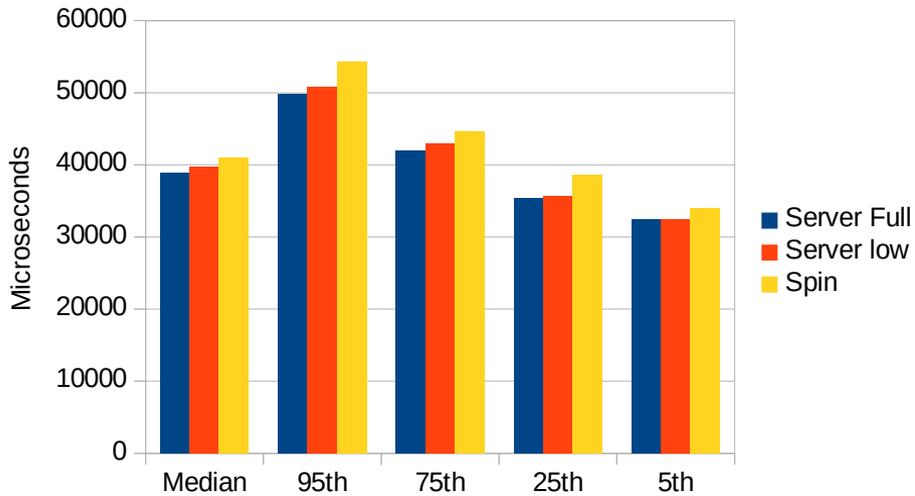| Sample size	
Server high	3609 (100%)
Server low	1088 (30%)
Spin	785 (22%)

303 samples rejected due to heuristics
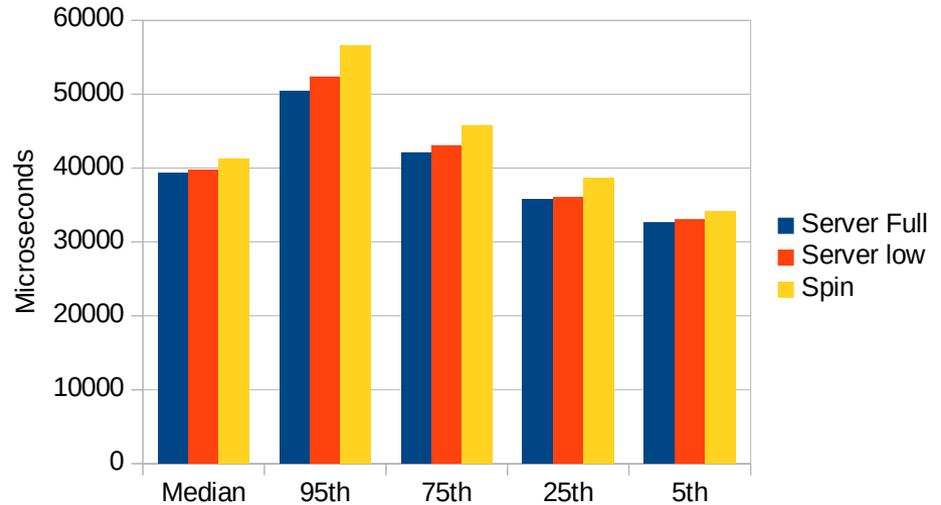
Sample size increases as reordering rate goes up.

Likely due to spurious loss detection and subsequent rate reduction (more RTTs to deliver the same amount of data)

RTT Estimates - loss

5% Random loss

10% Random loss

Number of RTT Samples

Uniformly distributed Random loss

5% loss

Sample size	
Server high	6761 (100%)
Server low	2018 (30%)
Spin	2017 (30%)

10% loss

Sample size	
Server high	8567 (100%)
Server low	2785 (32%)
Spin	2784 (32%)

Sample size increases as loss rate goes up.
Due to loss detection and subsequent rate reduction (more RTTs to deliver the same amount of data).