SAND2017-9799C

Exceptional service in the national interest

Sandia National Laboratories

Measuring Minimum Switch Port Metric Retrieval Time and Impact for Multi-Layer InfiniBand Fabrics

Michael Aguilar

Benjamin Allan and Serge Polevitzky

Sandia National Laboratories

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Outline

- Motivation
- Technical goals
- Experimental Approach
- Results
- Conclusions and Preliminary Recommendations
- Future Work

Motivation

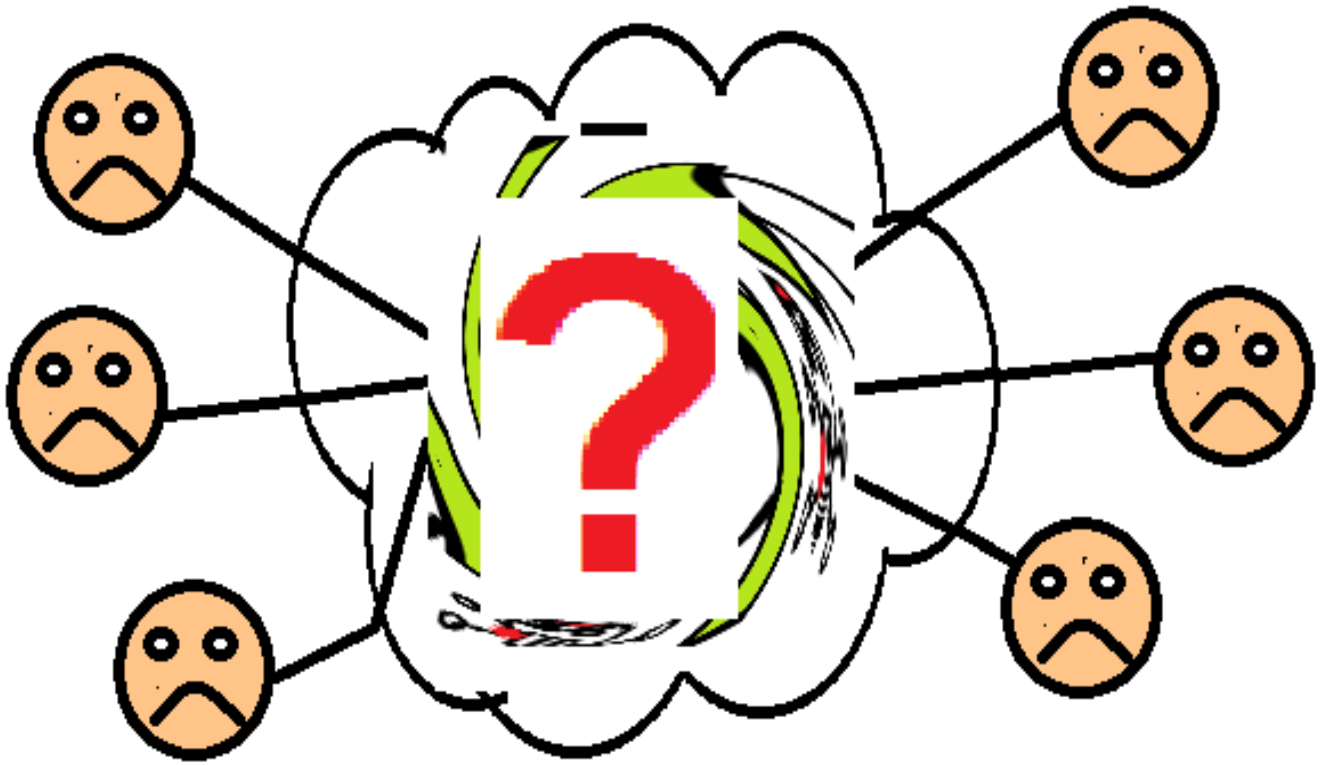

Congestion

Inconsistent run-time
Bad application?
The cluster failed?

Motivation

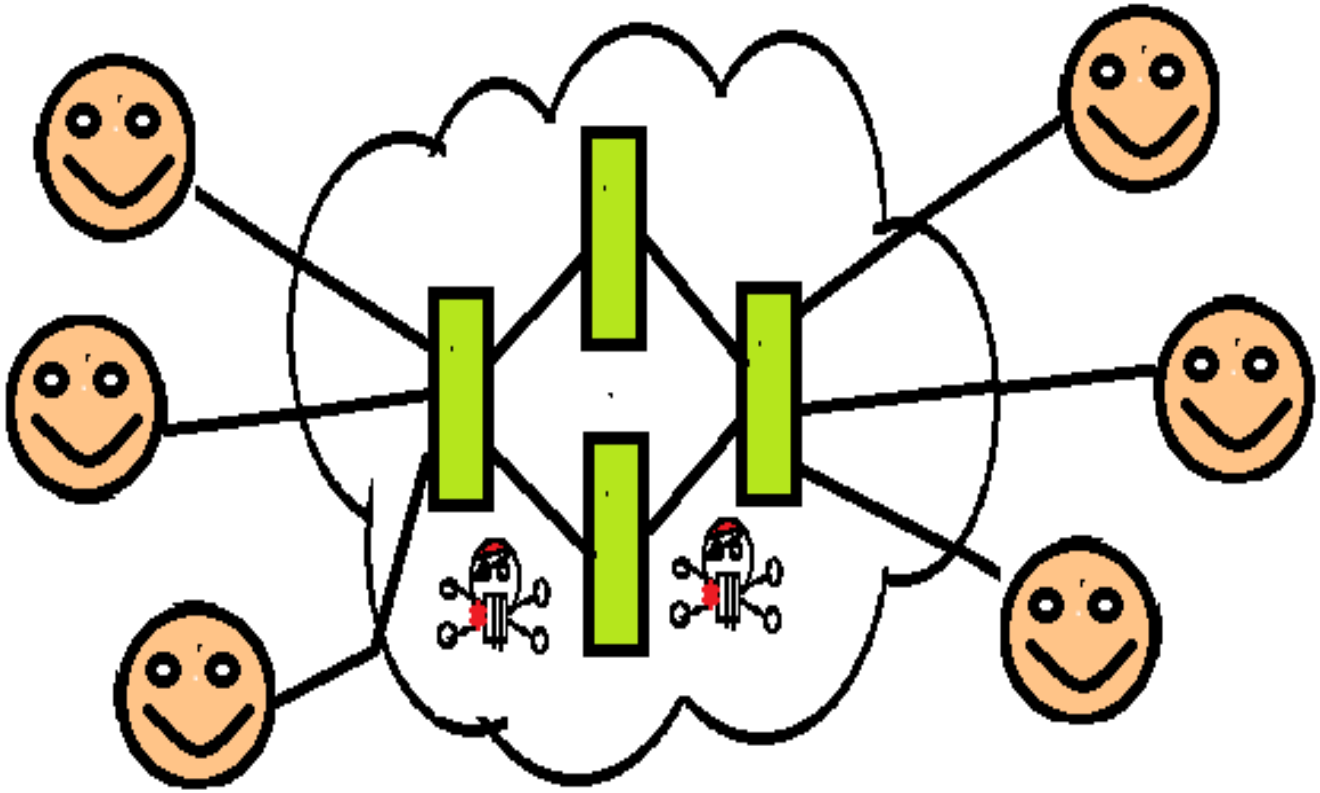

Congestion

Inconsistent run-time
Bad application?
The cluster failed?

Motivation

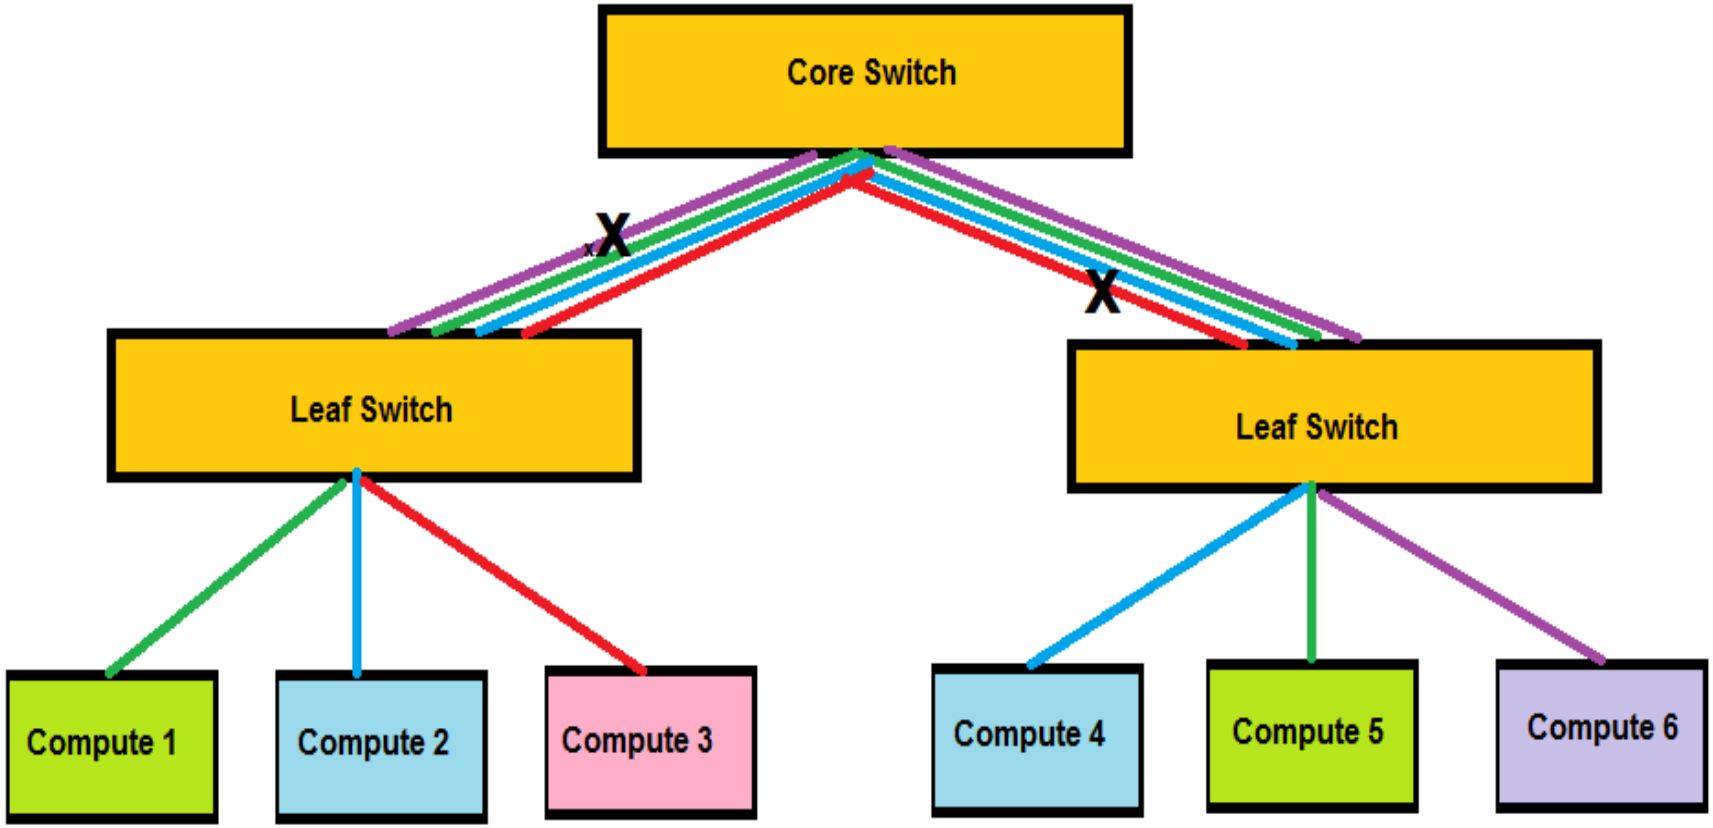
Core Switch
X
X
Leaf Switch
Leaf Switch
Compute 1
Compute 2
Compute 3
Compute 4
Compute 5
Compute 6

Motivation

- Identify network problems in initial stages
  - Requires high fidelity monitoring of many network parameters
    - E.g., symbol errors, congestion, utilized bandwidth
  - Enables proactive mitigation and user notification
- Understanding impact of network congestion on application performance
  - Requires knowledge of congestion levels
- Understanding network loads
  - Requires knowledge of traffic loads and patterns
- Planning future system network provisioning
  - Requires both of the above

Query Performance Concerns

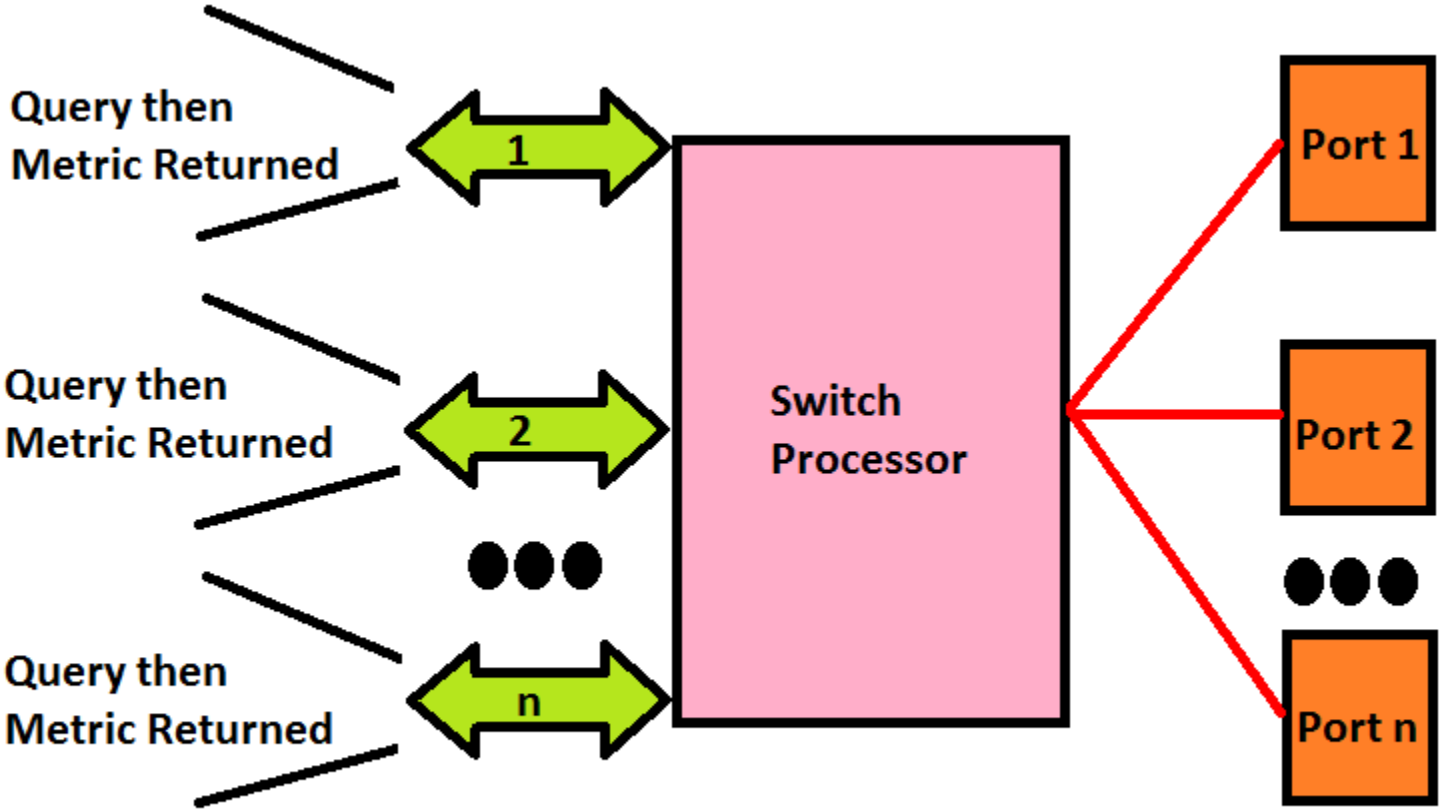

Time Spread

What is the minimum amount of time necessary to get data from a set of InfiniBand switches?

- The time spread between first port sample acquisition and last port sample acquisition.

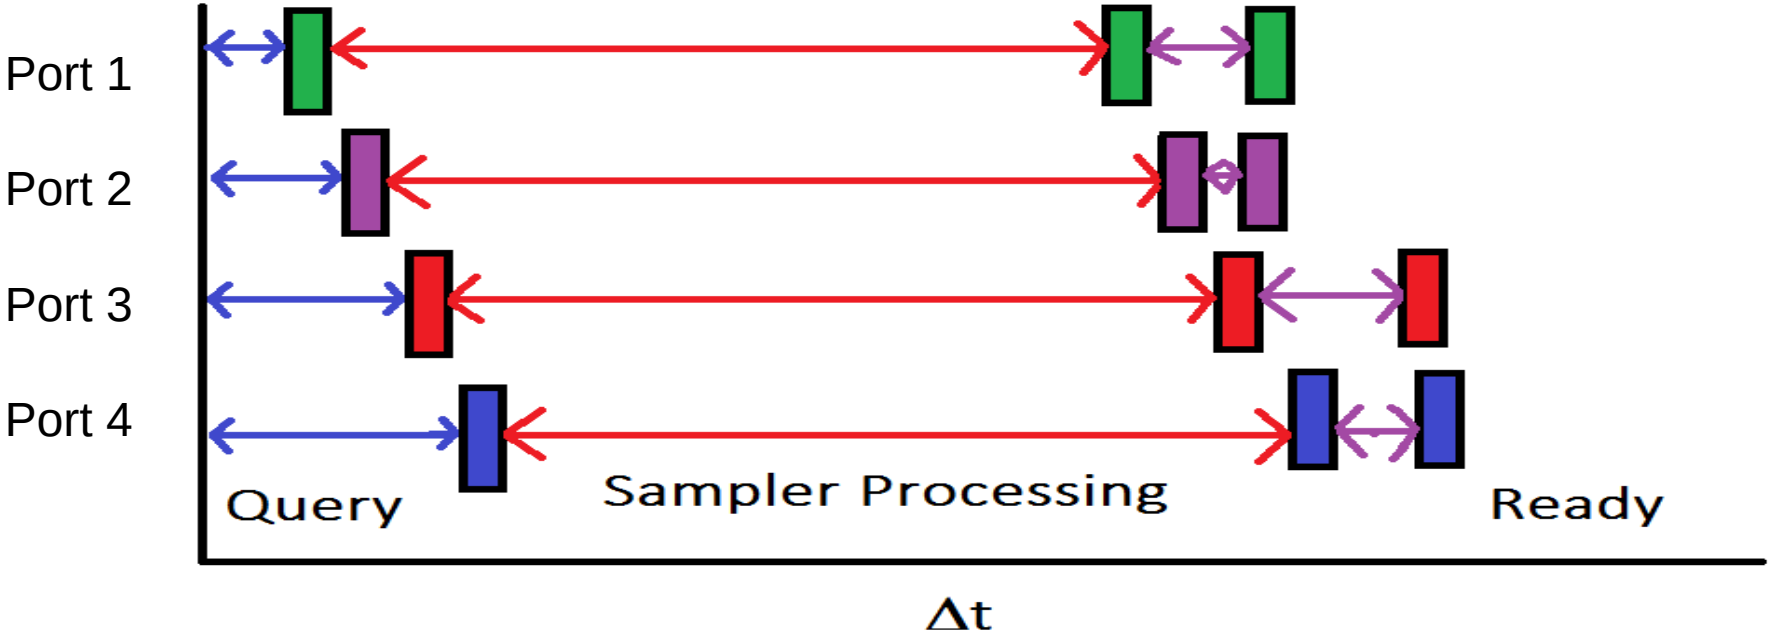

Experimental Approach

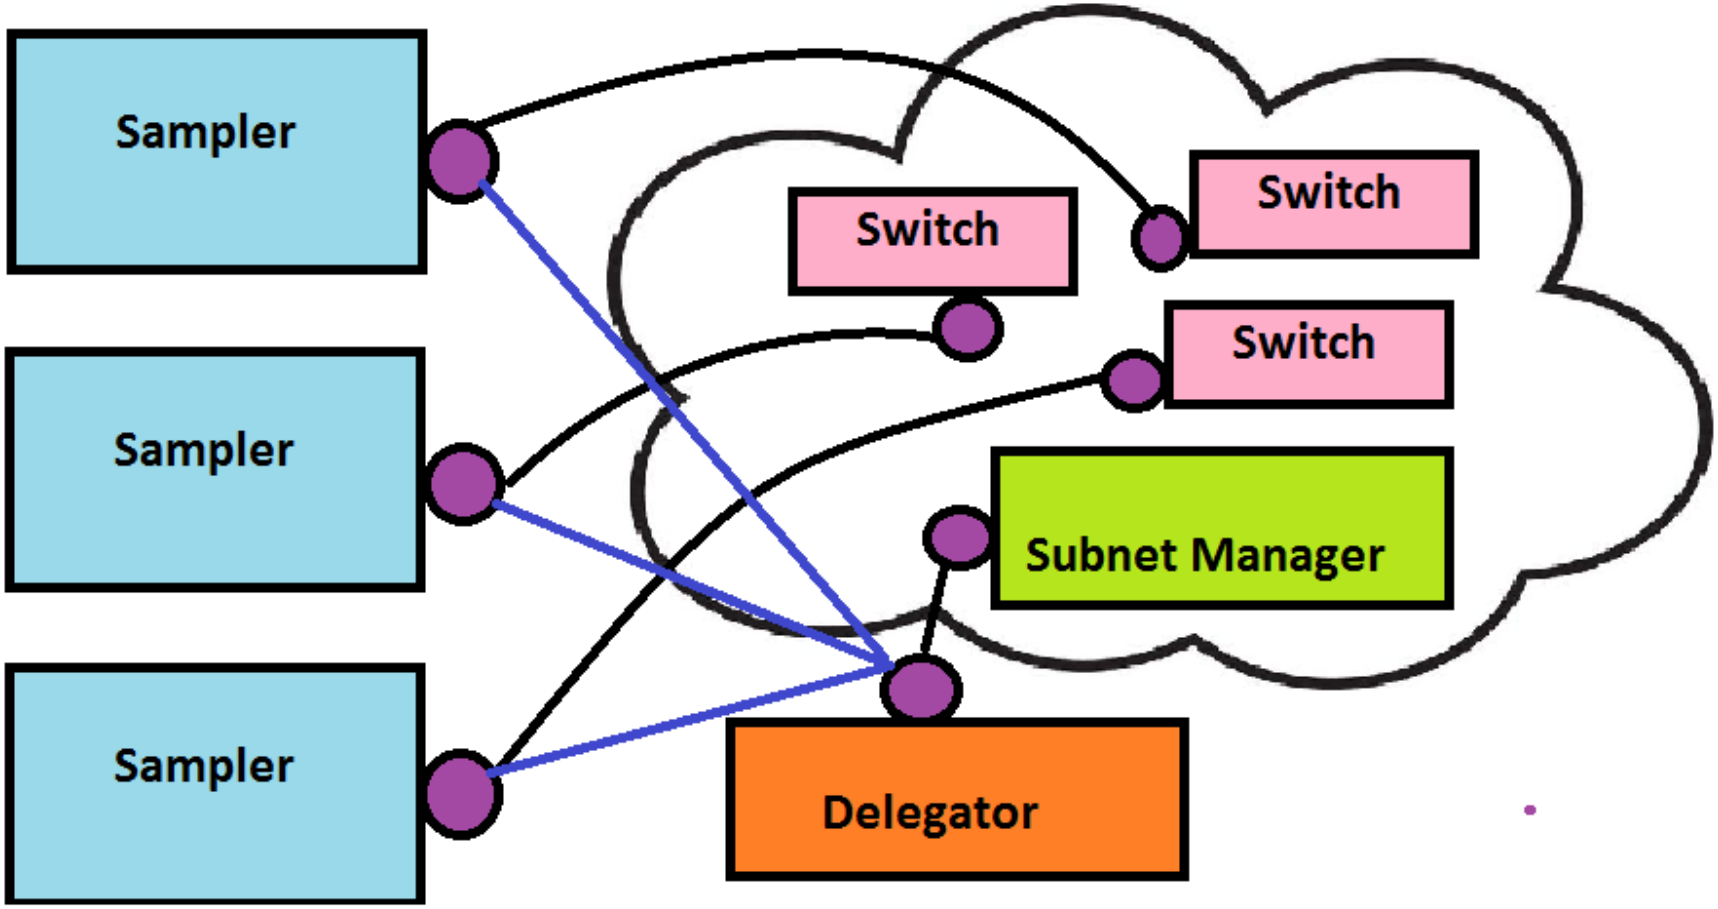

Questions of Concern

- What is the amount of interference by IB queries with application network bandwidth?
- How much do applications interfere with IB data collection?
- What is the minimum spread time for port data collection?
- What are the worst case effects of redundant port data collection?
- What is the impact of topology (hops, multiple targets) on spread?

Experimental Approach

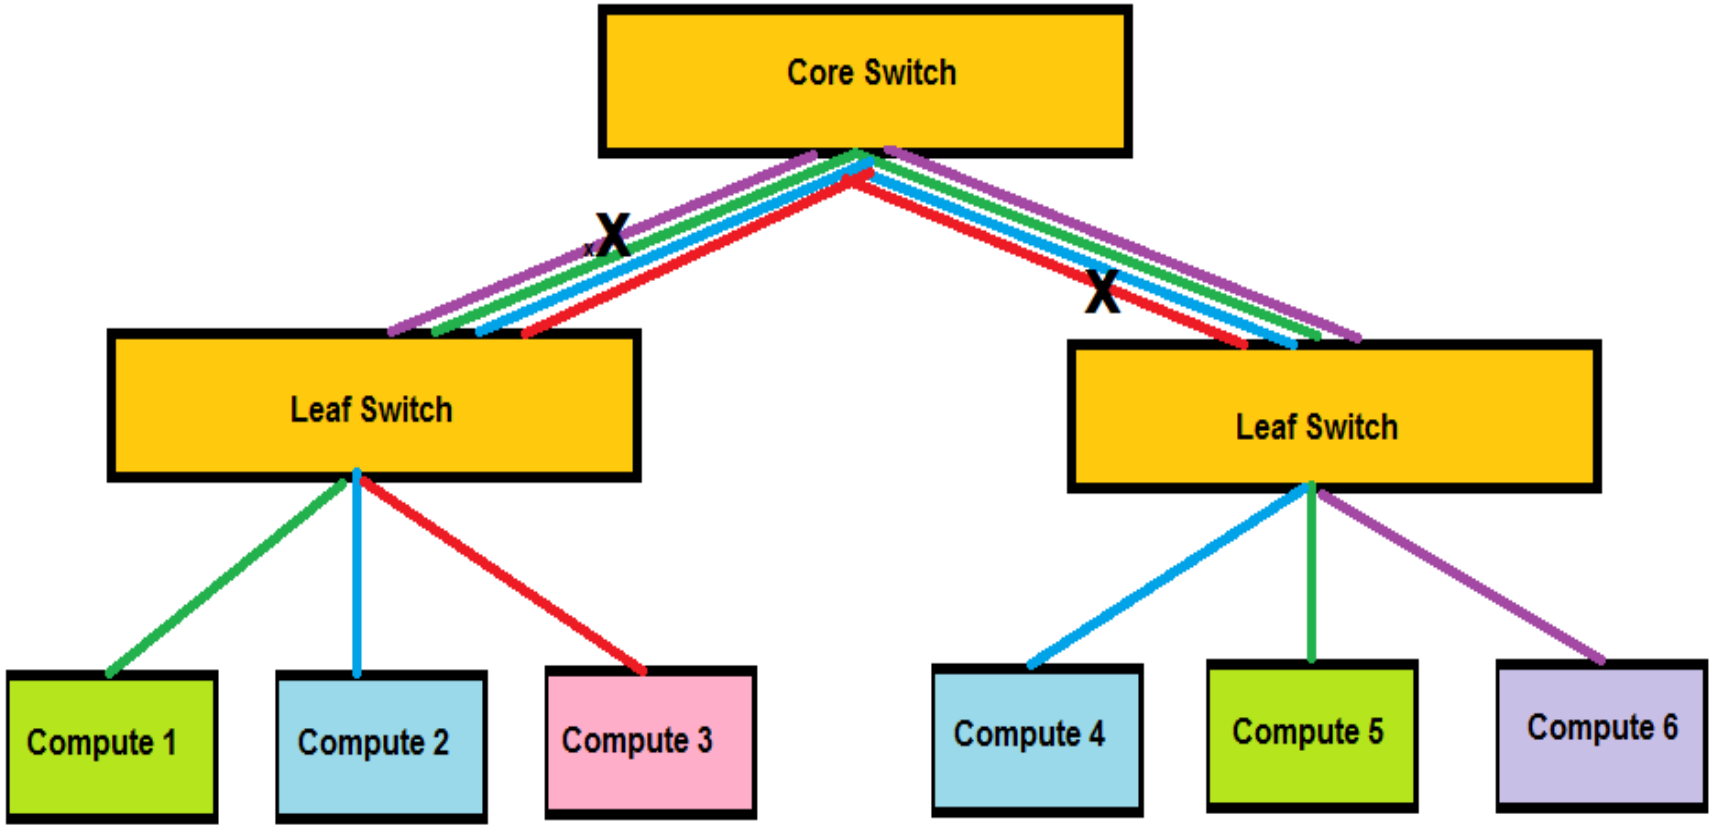

What is the amount of interference by IB queries with application network bandwidth?

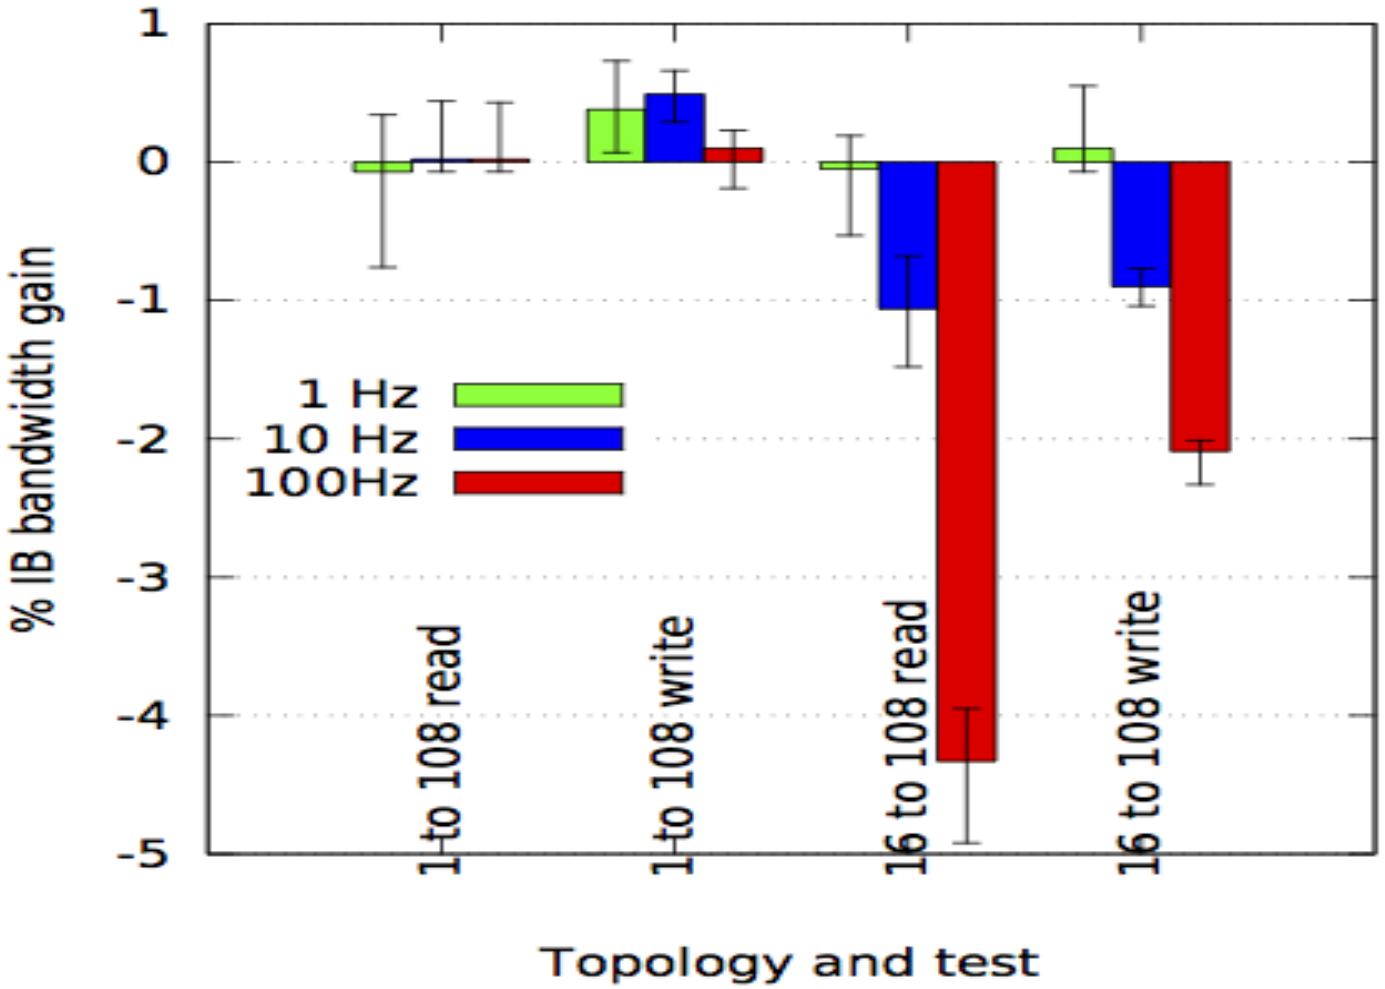

How much do applications interfere with IB data collection?

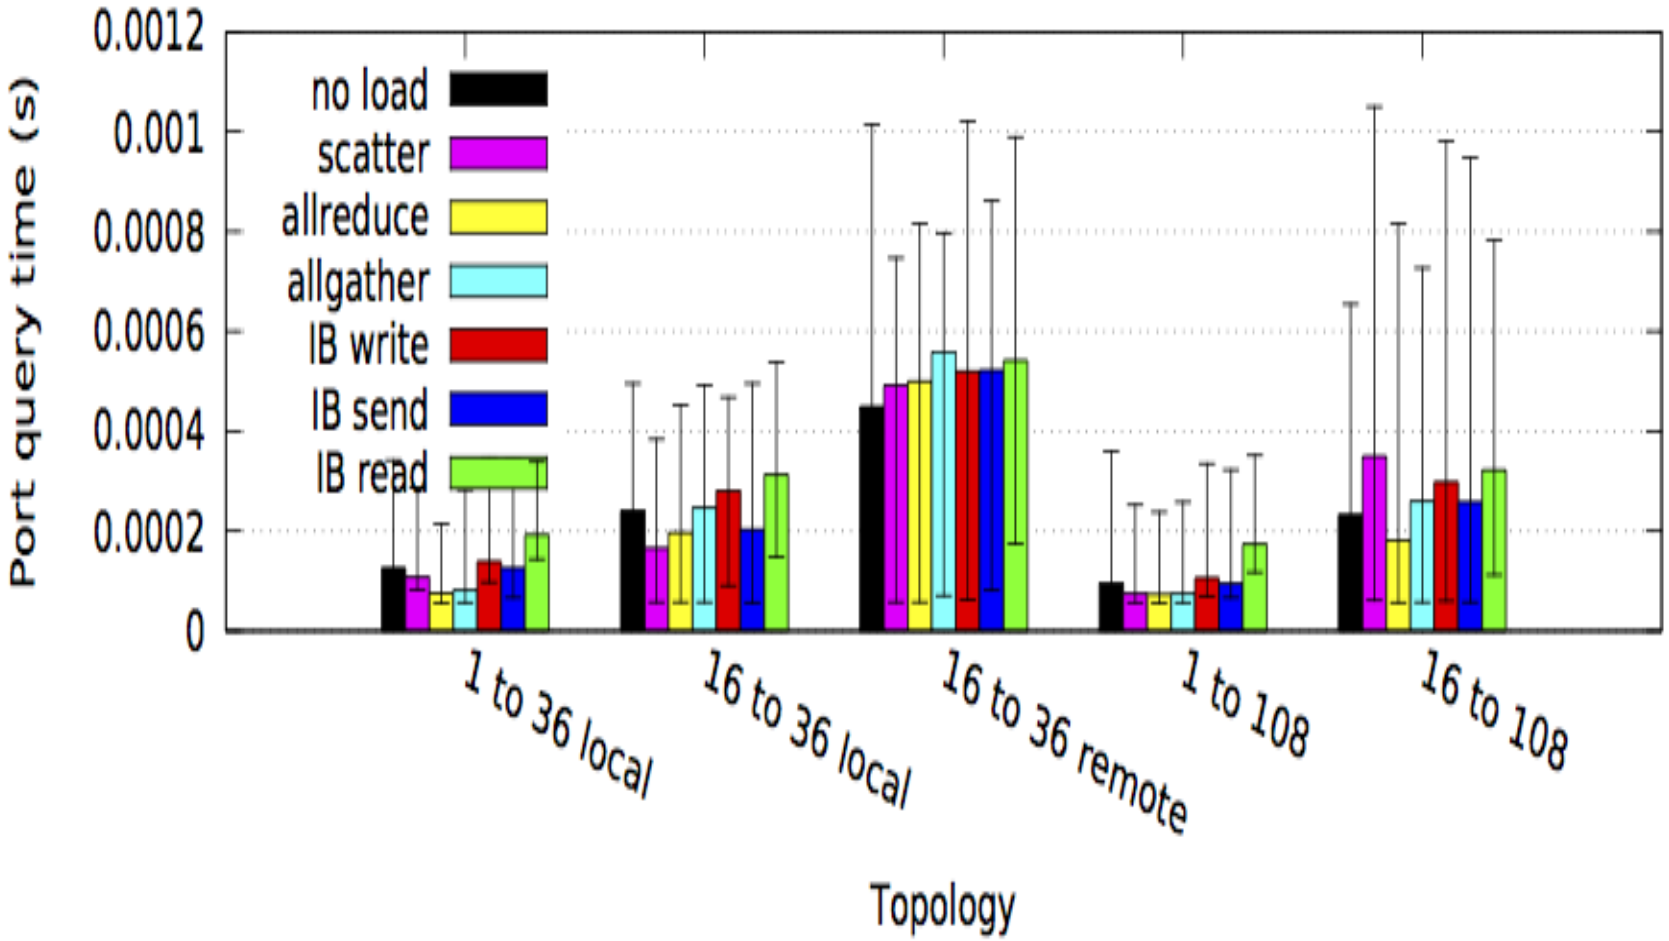

What is the minimum spread time for port data collection?

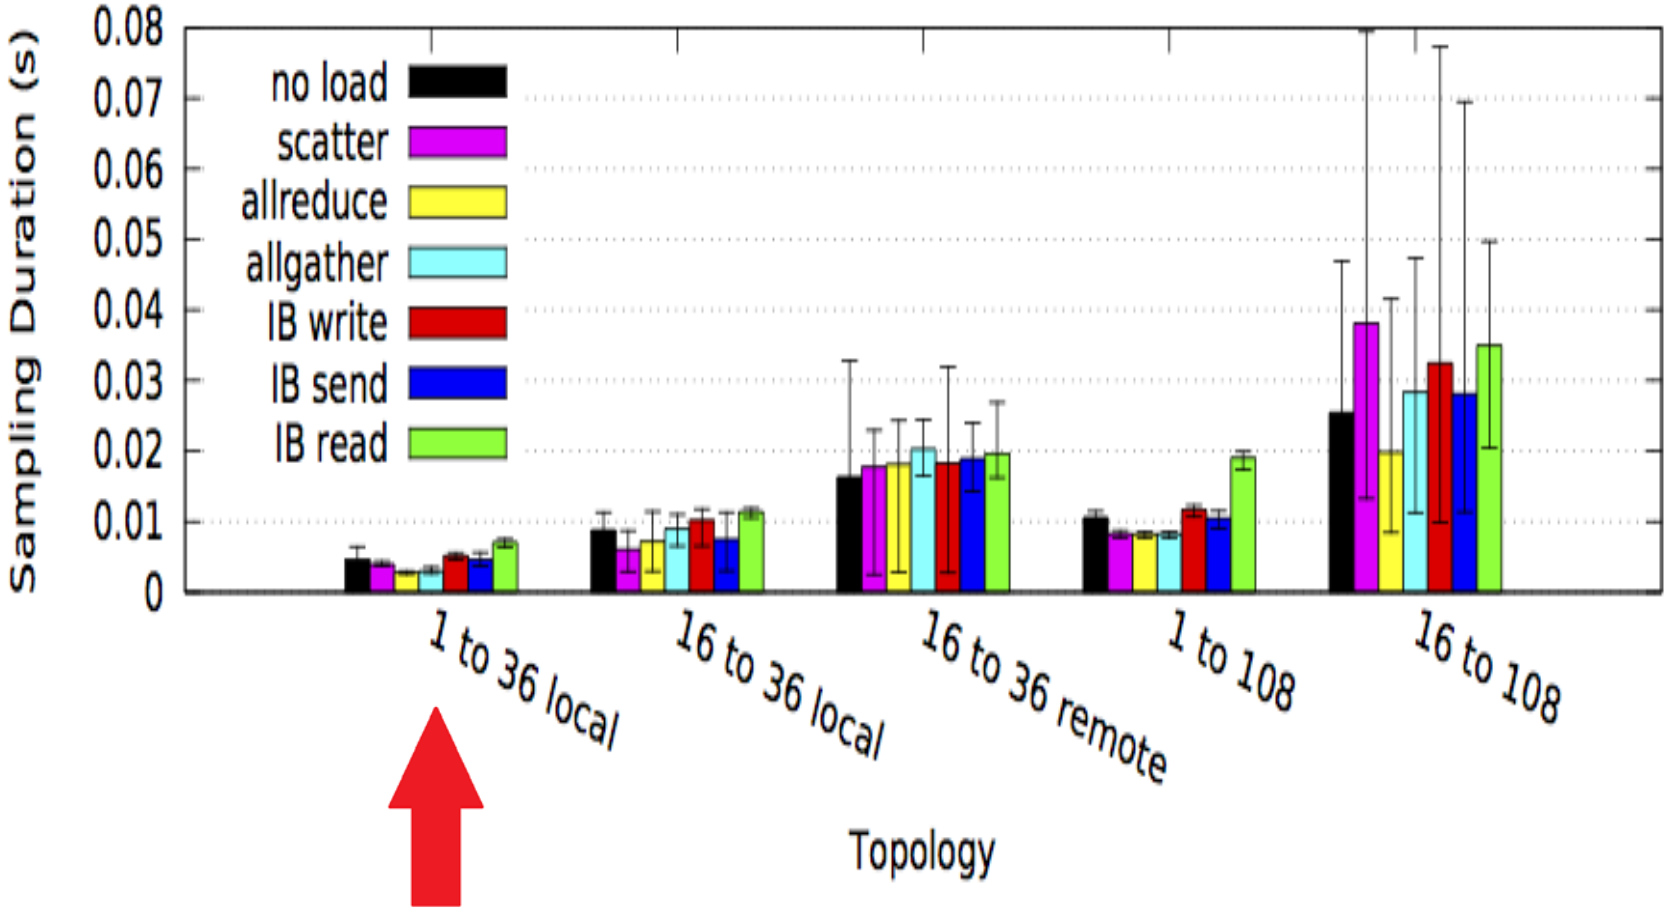

What are the worst case effects of redundant port data collection?

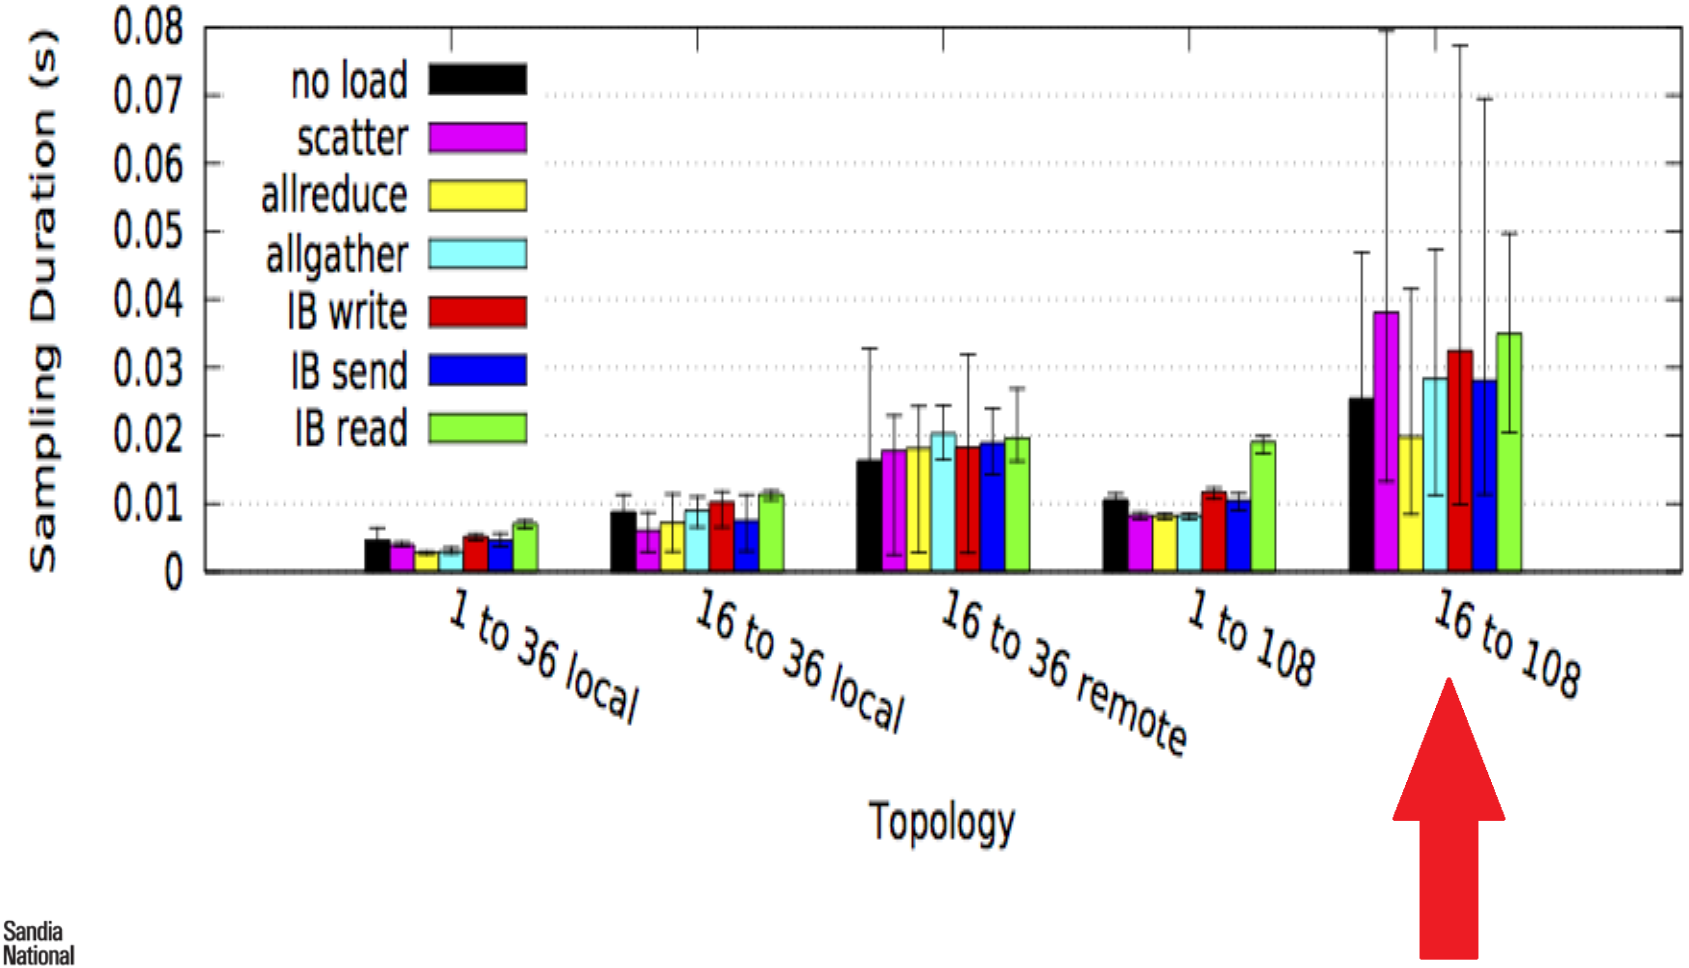

What is the impact of topology (hops, multiple targets) on spread?

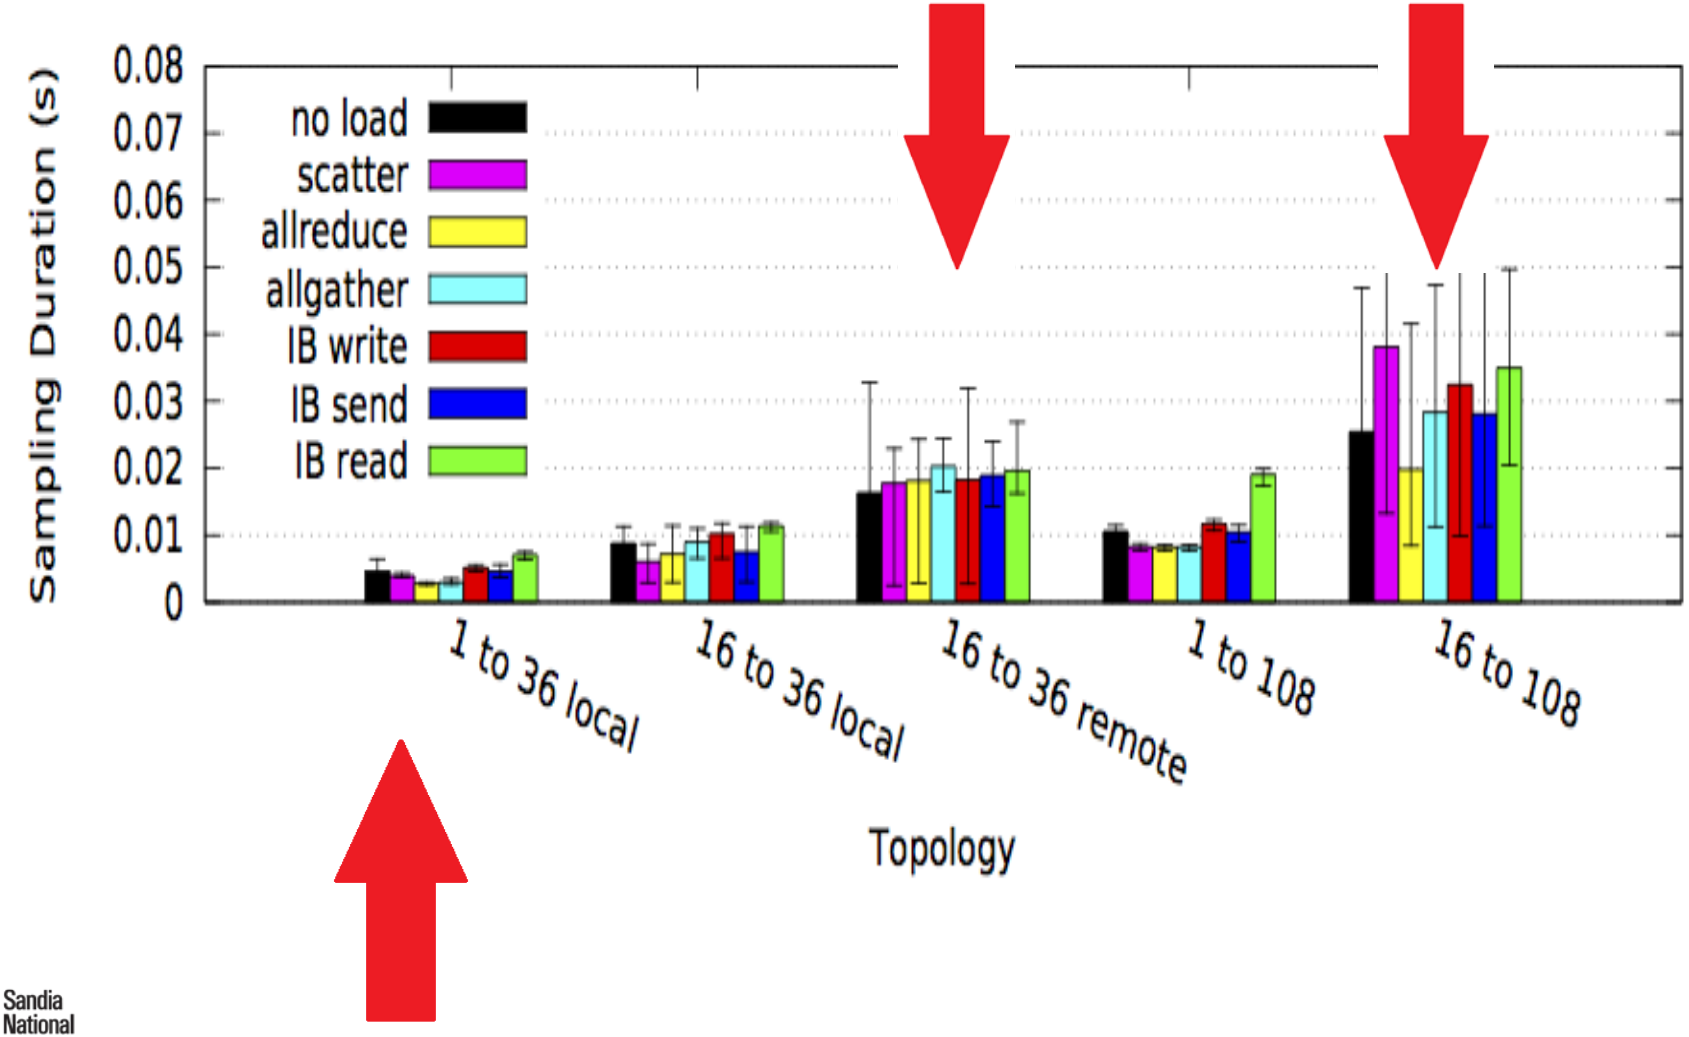

Initial Conclusions and Preliminary Recommendations

- Parallelism in collection efforts, where the sampler is local to a switch, reduces query times.
- We can measure IB switch performance metrics with negligible bandwidth overhead at frequencies up to 100 Hz if the topology is correctly laid out.
- A slight improvement in application traffic performance has been seen at sampling rates of 1 Hz and 10 Hz.

Further Work

- Continue experiments on a larger HPC system with a larger fabric.
- Distributed metric gathering
- Determine the optimum sampler topology
- Publish the ibfabric sampler plugin so that others can use it.

Questions?