
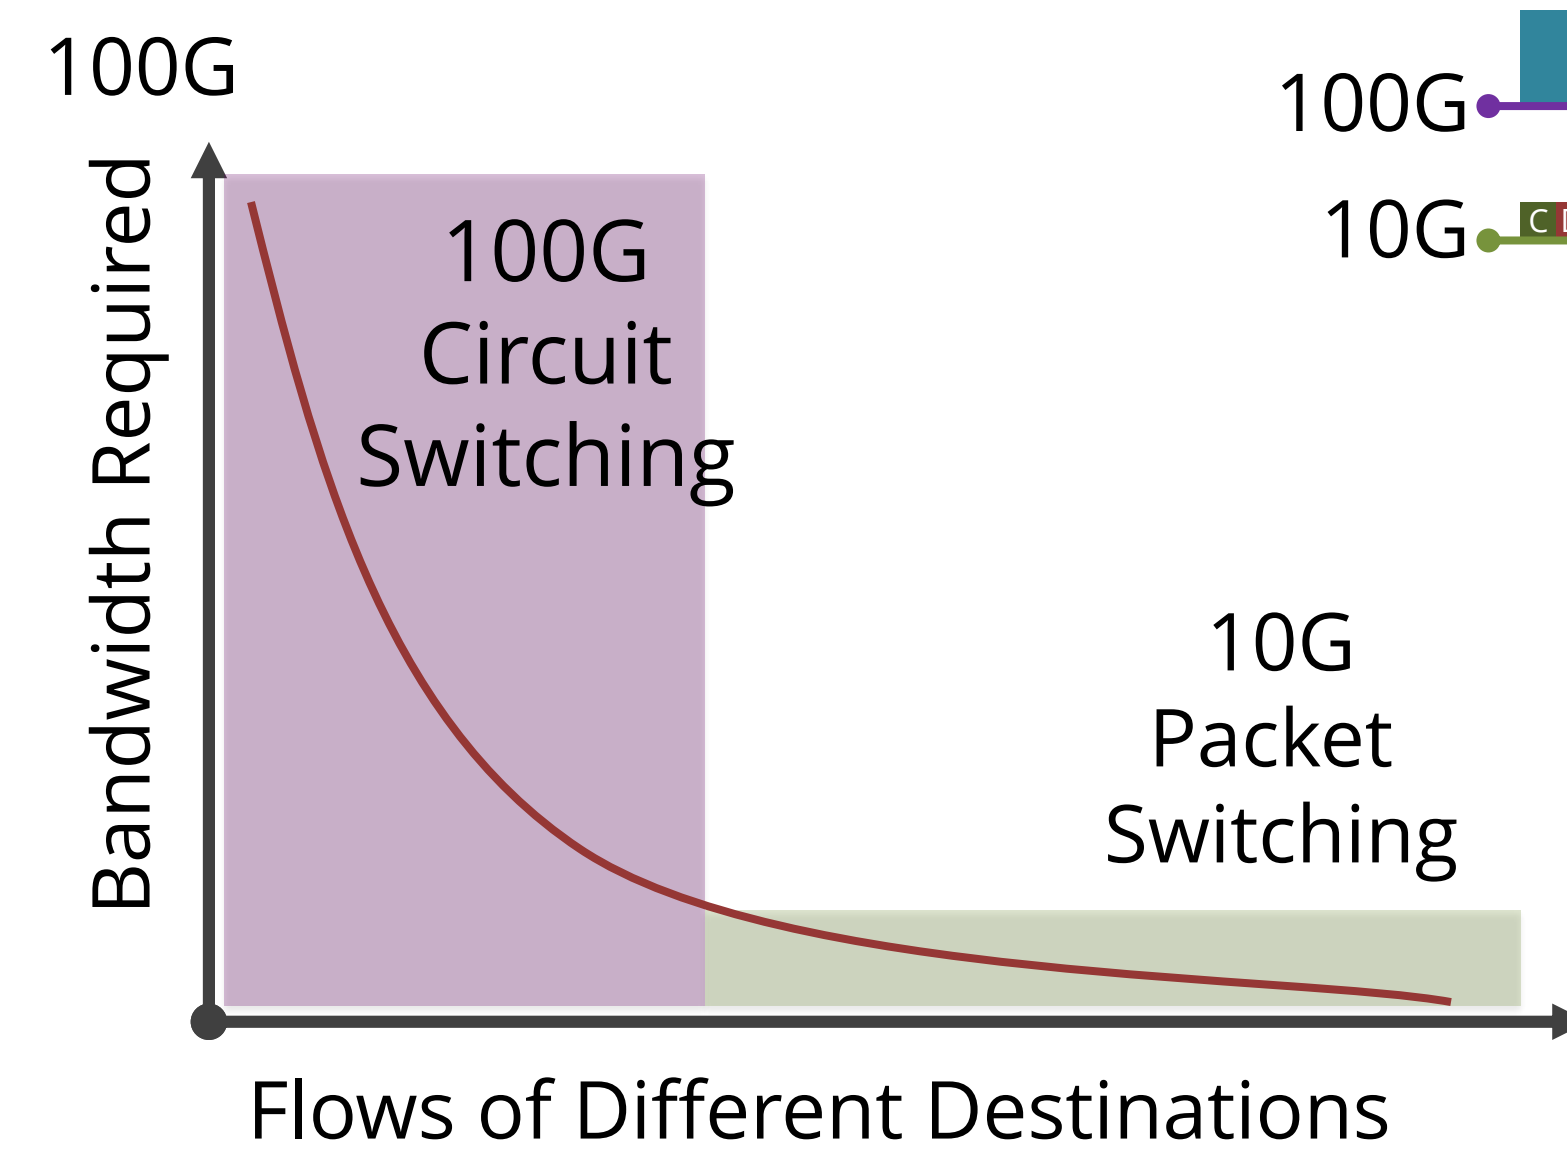


Switch Scheduling in a Hybrid World

He Liu, George Porter, George Papen, Stefan Savage, Geoffrey M. Voelker, Alex C. Snoeren


REACToR: 100G Hybrid Datacenter Switches



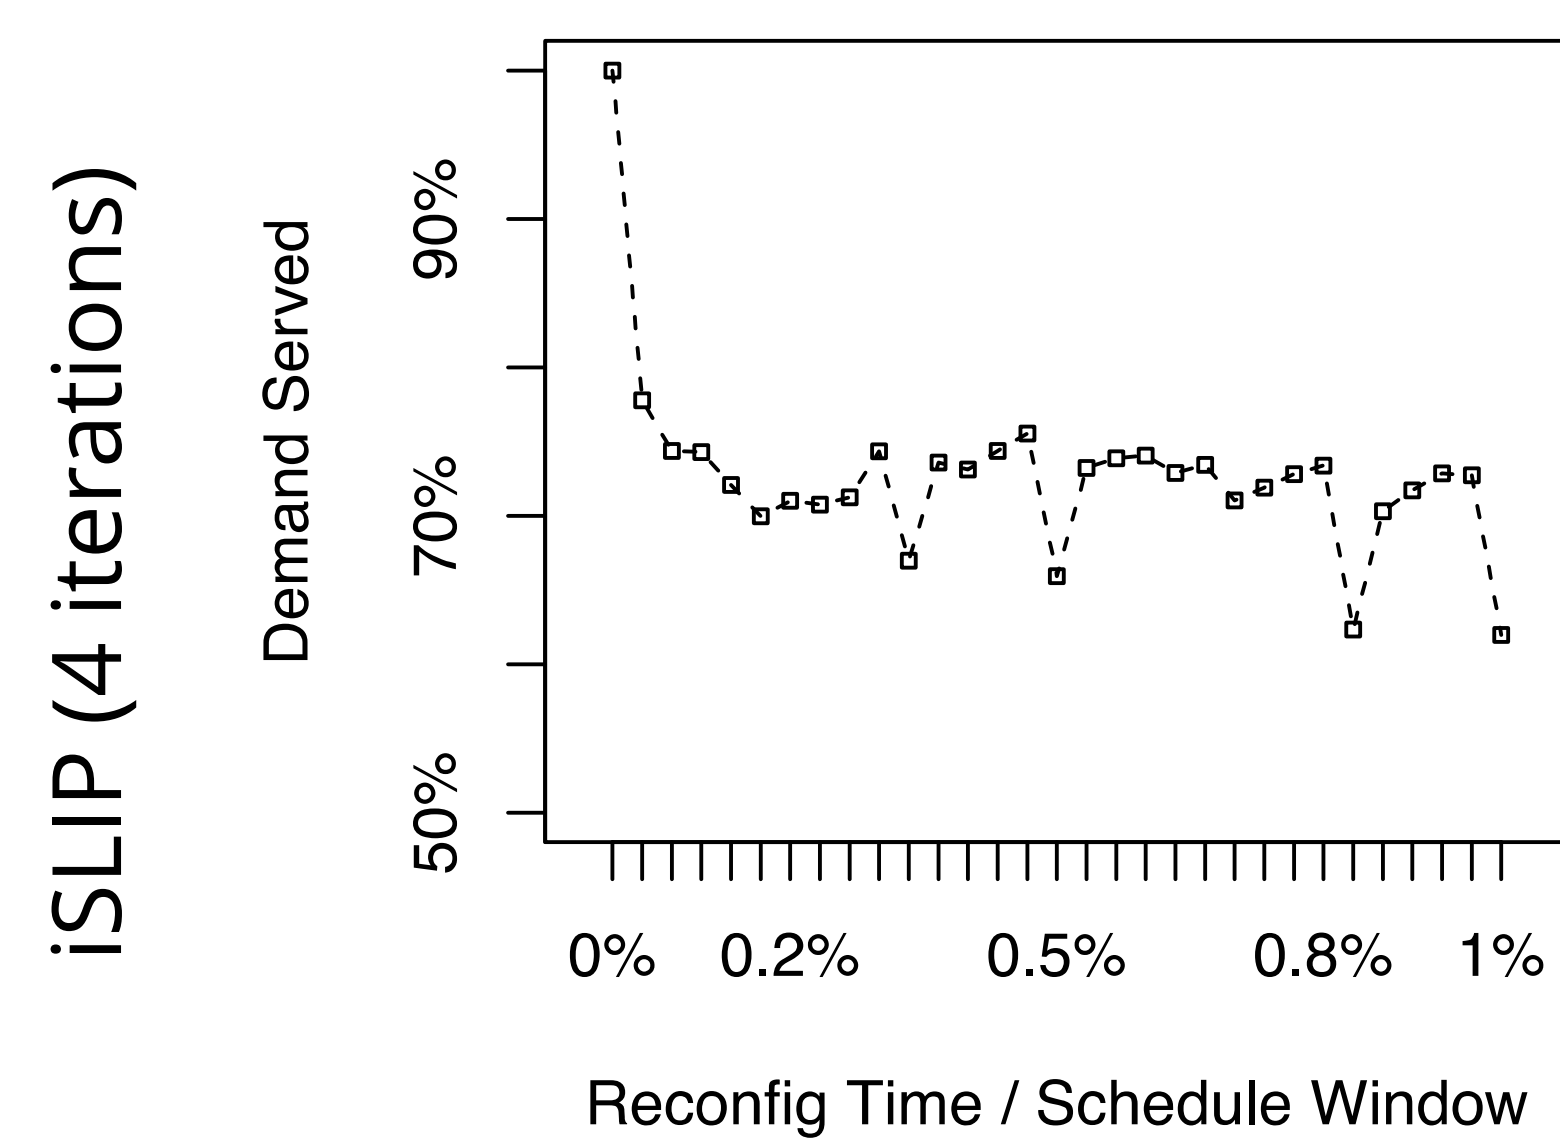
- Performance akin to electrical packet switch
 - Fewer Transceivers
 - Lower cost
- [NSDI 2014]

TDMA Requires a Scheduler

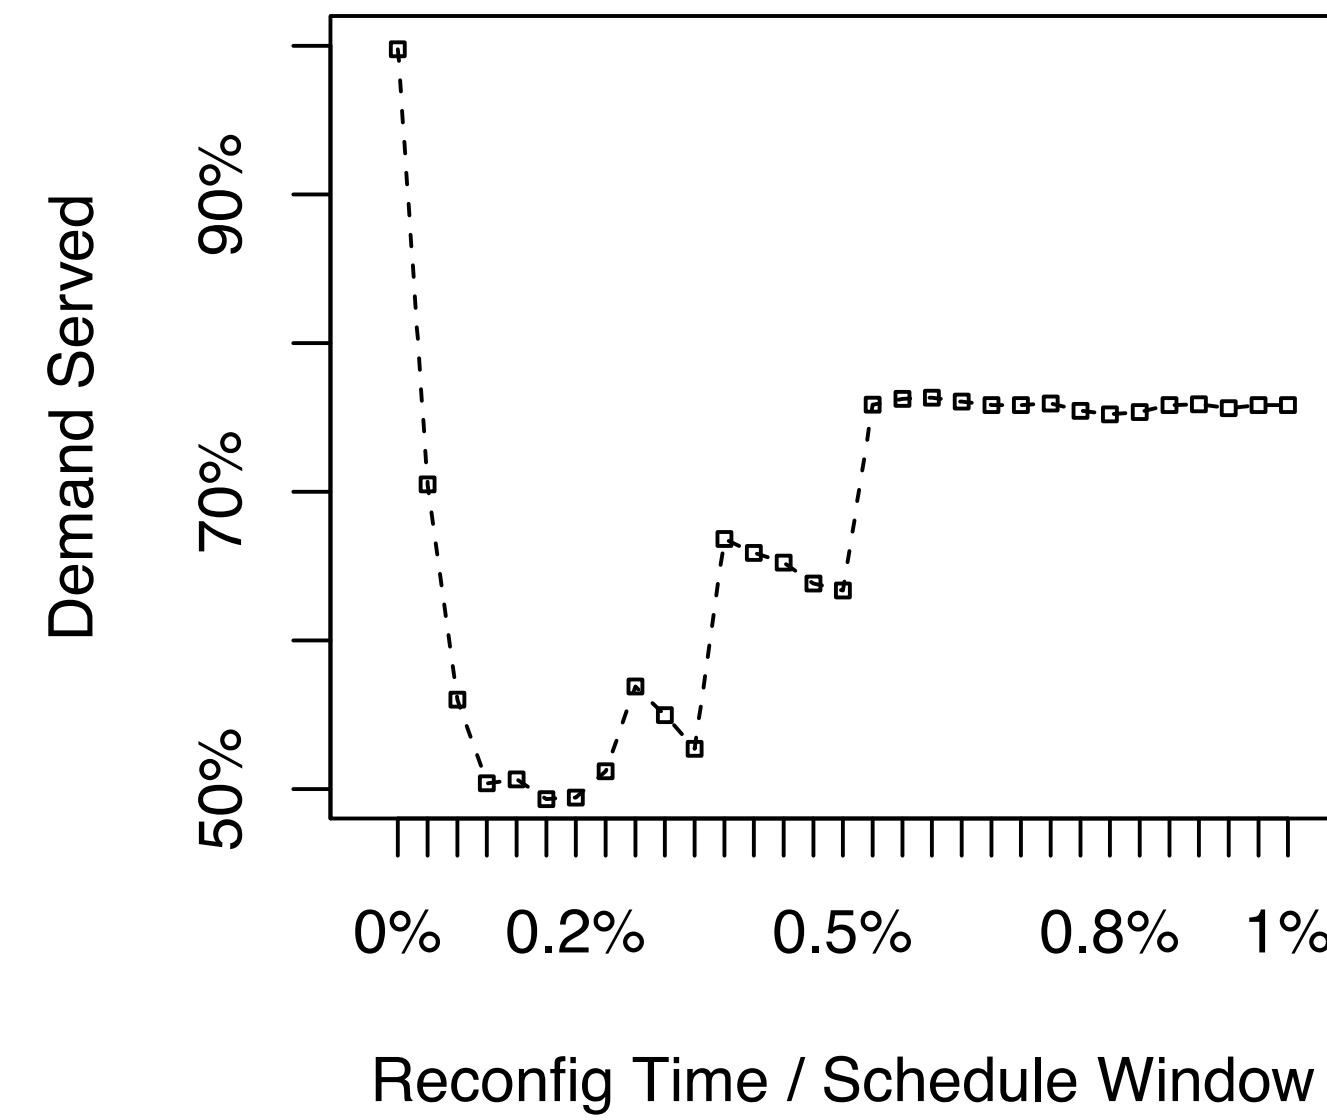
- Reconfigures in ~10 microseconds; TCP at line rate
- Simple Example: round-robin schedule
- Skewed demand in real life; needs more complicated algorithm



Previous Scheduling Algorithms Don't Perform Well



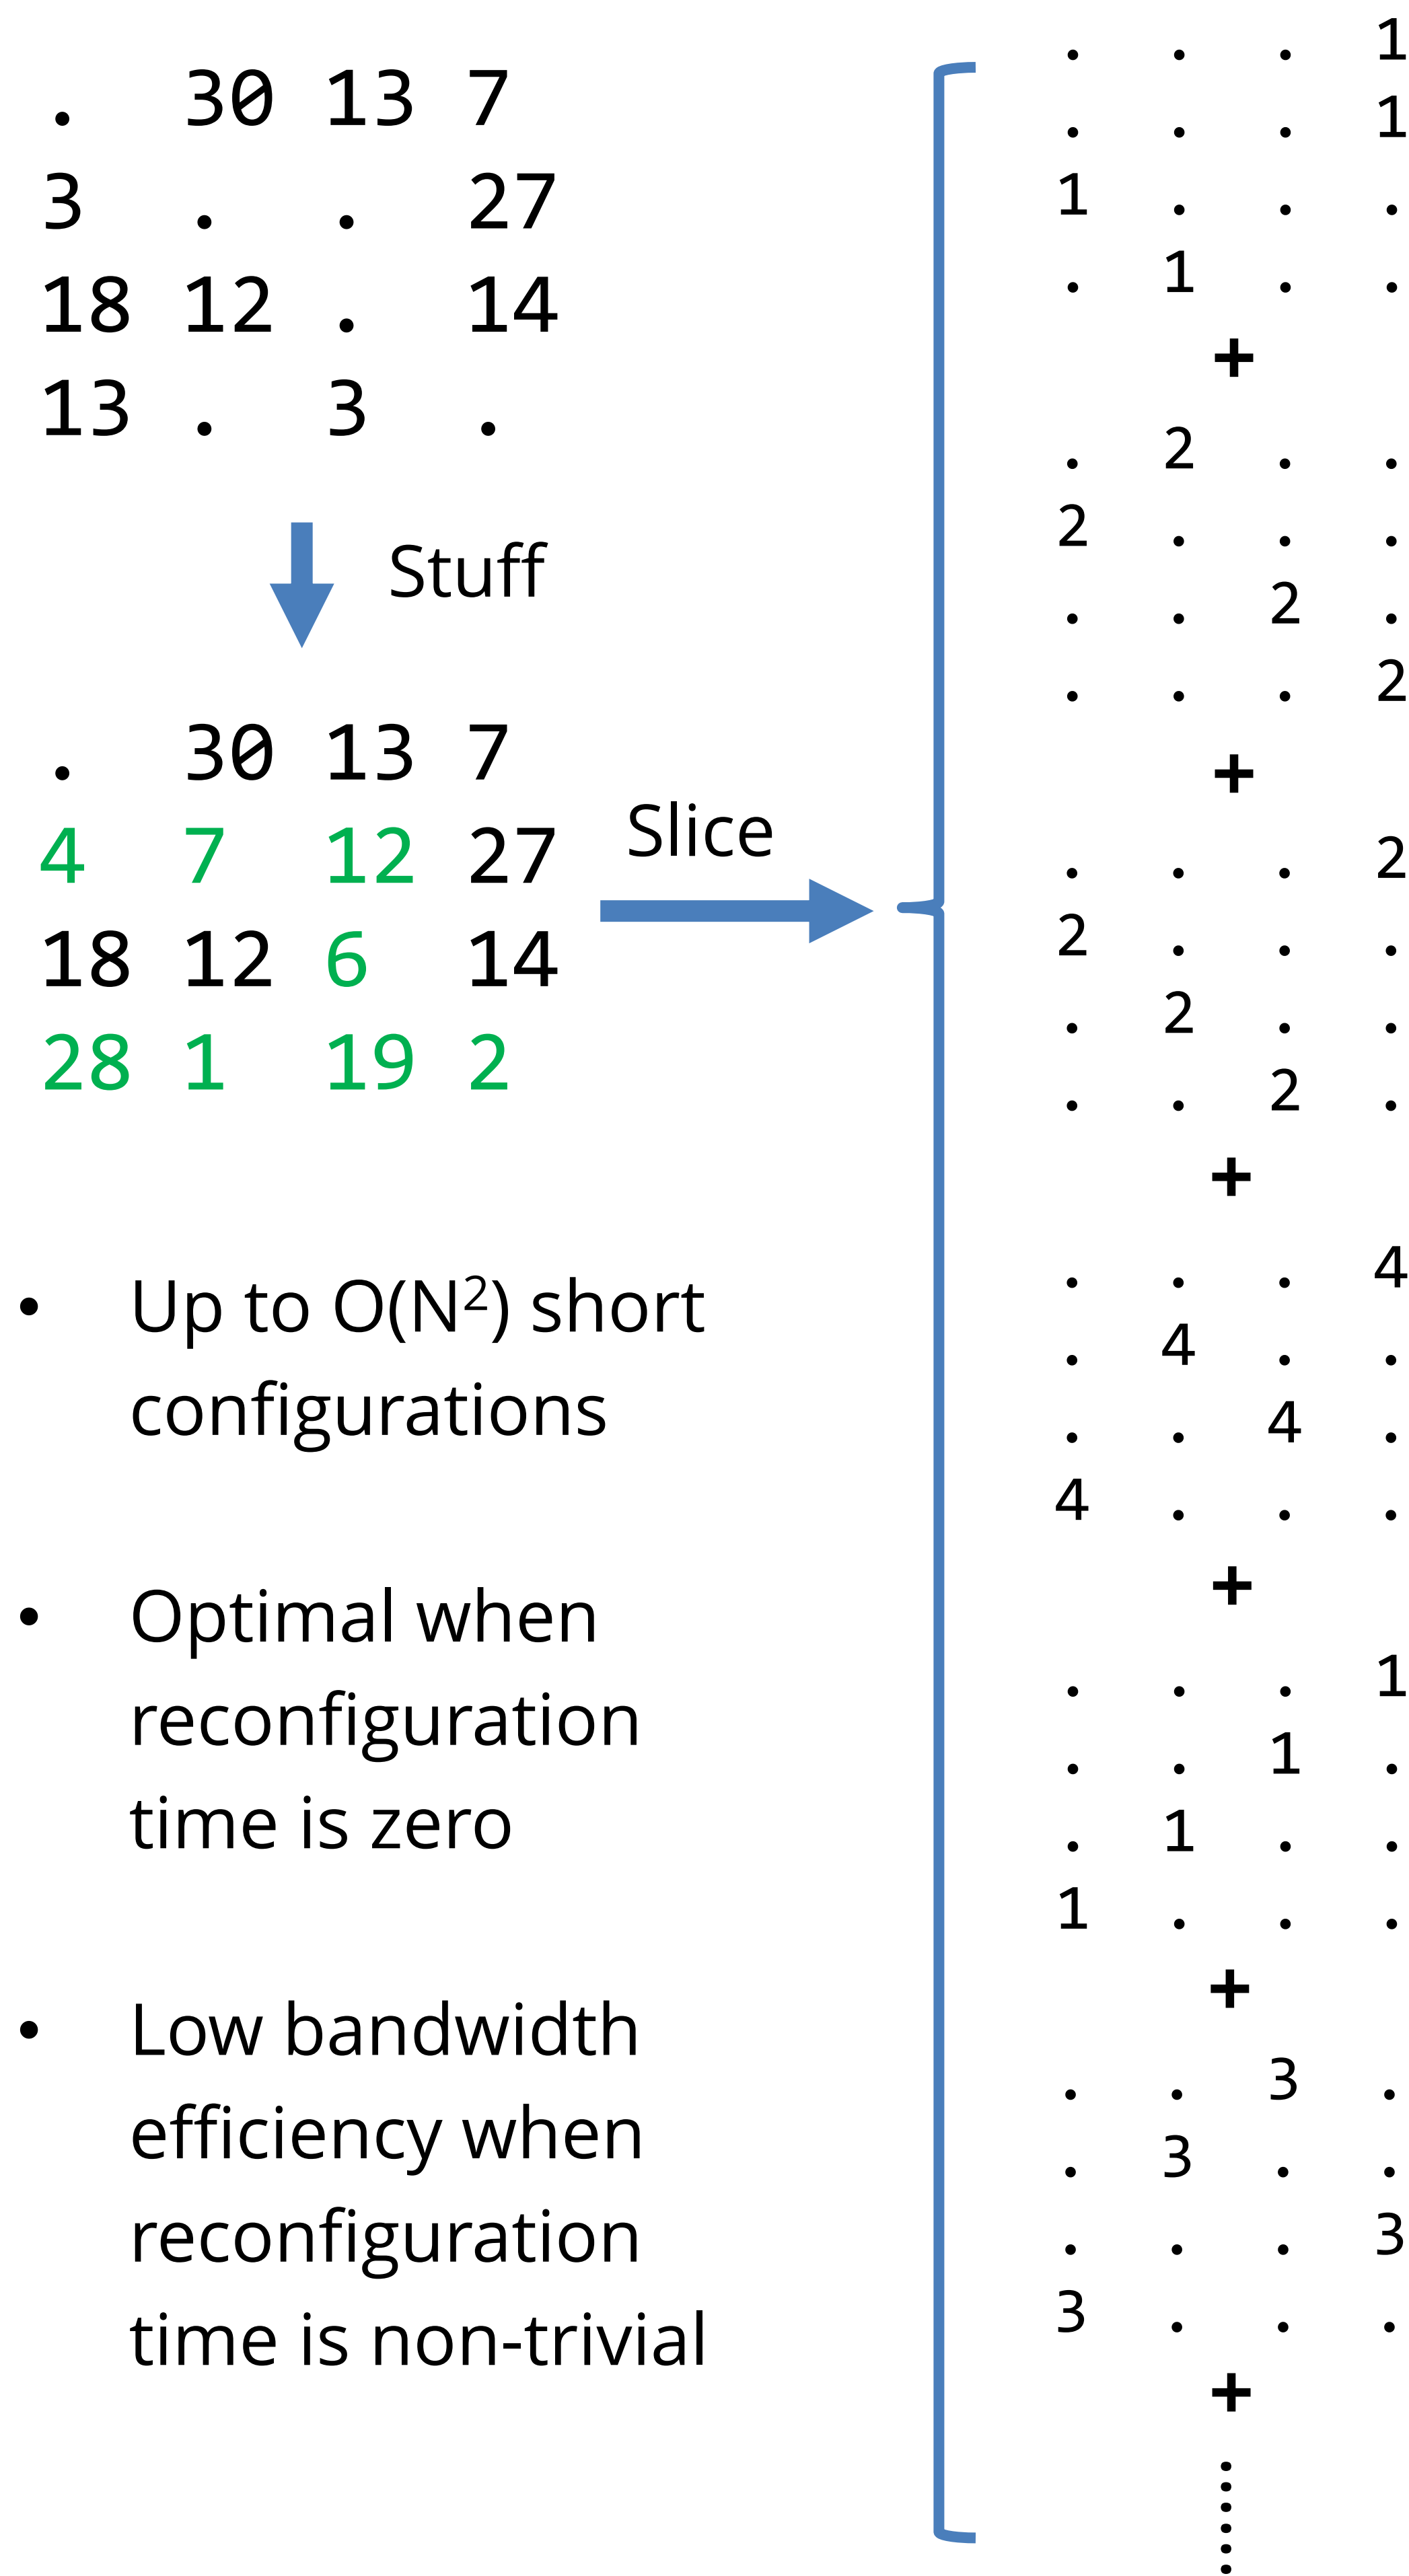
Birkhoff-von Neumann Decomposition



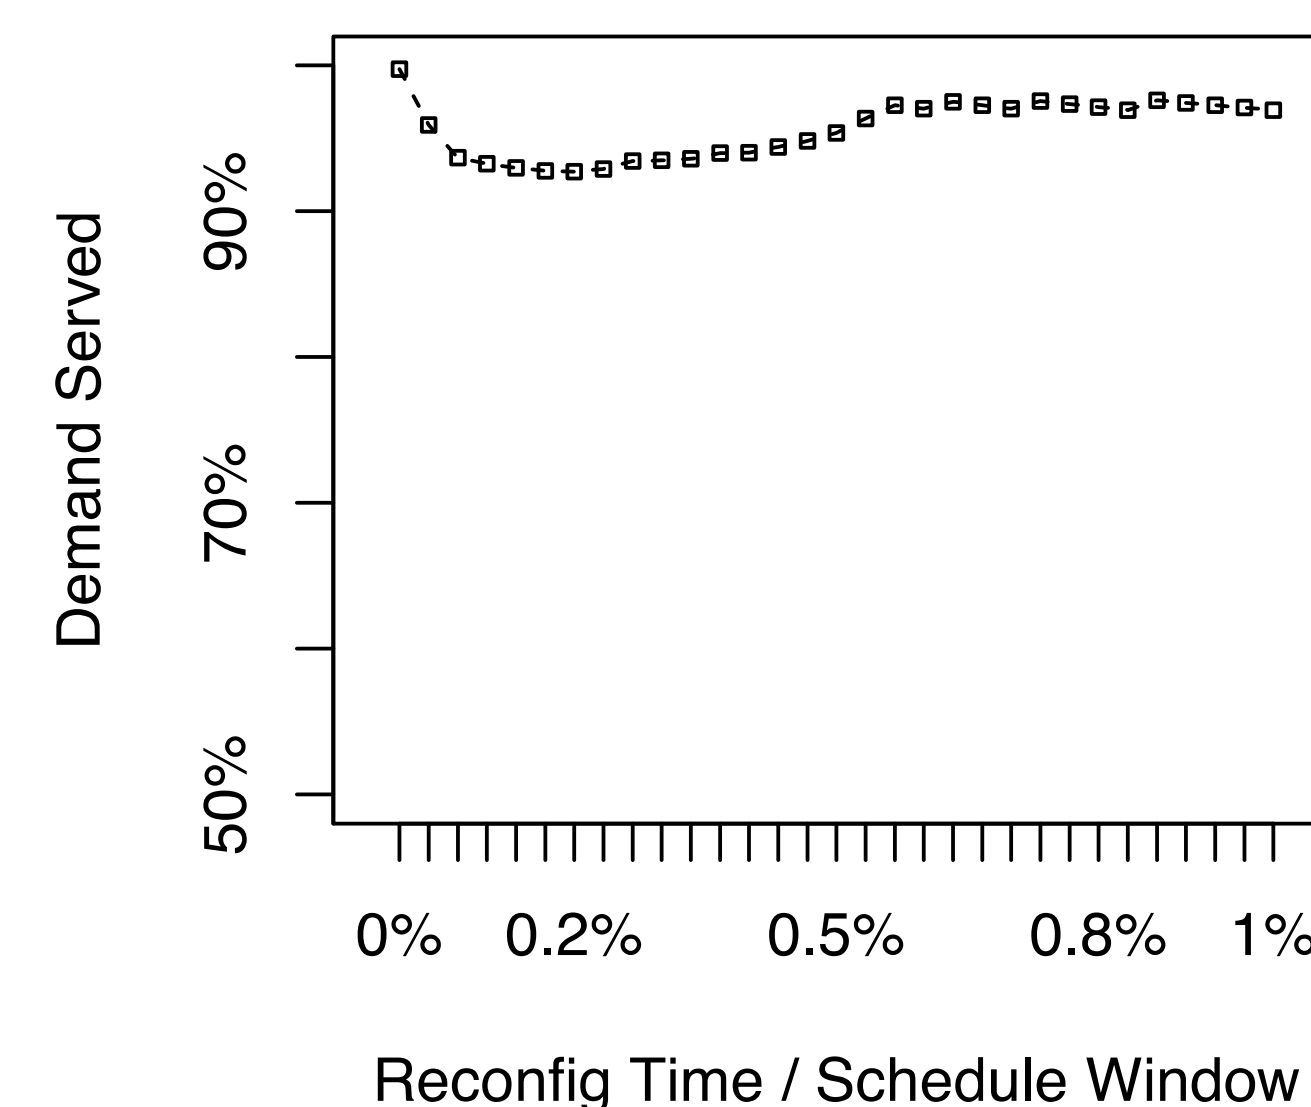
Reasons:

- Optical switching; no buffering in the switch; no speedup
- Incur a reconfiguration delay on every change

How Conventional BvN Works

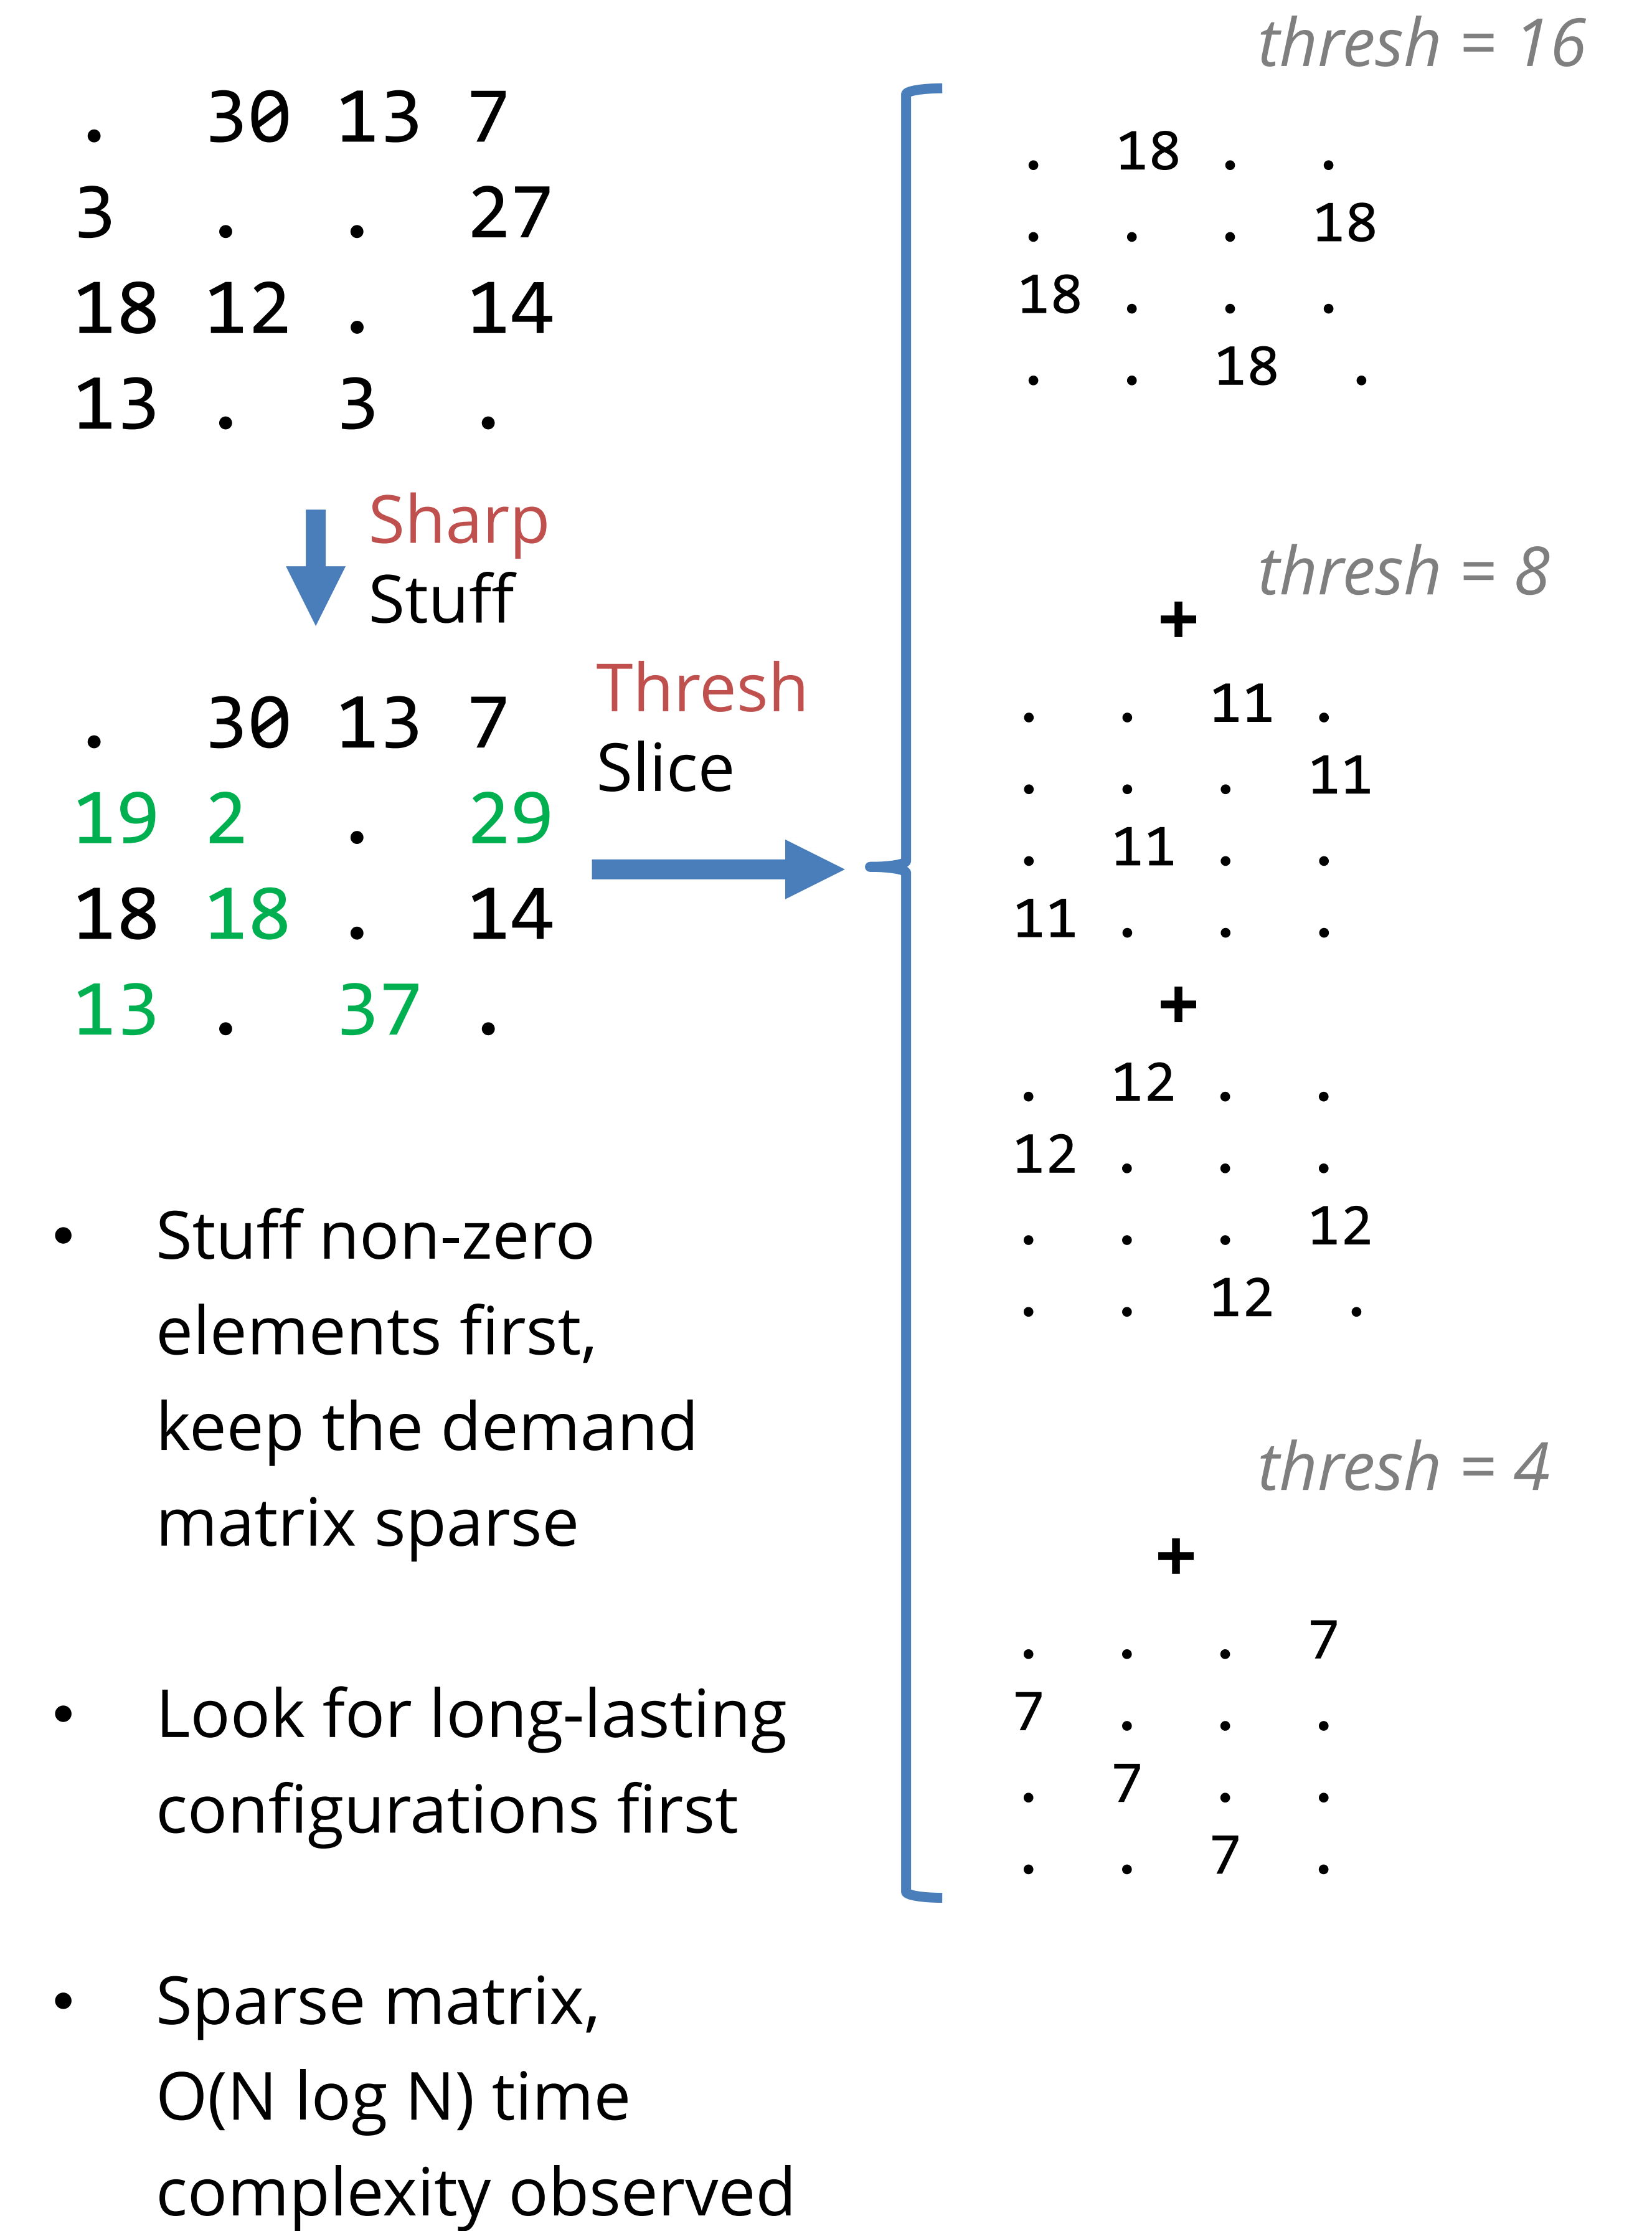


Performance Evaluation



- With hybrid setup, throughput akin to 100G electrical packet switch
- $O(N \log N)$ time complexity observed for sparse input matrices

How Solstice Works



- Stuff non-zero elements first, keep the demand matrix sparse
- Look for long-lasting configurations first
- Sparse matrix, $O(N \log N)$ time complexity observed

Ongoing Work

- Integrate with demand estimation
- Proof of time complexity
- Hardware implementation
- Distributed solution

