Cornell University

Data Center Networks and Switching and Queueing and Covert Timing Channels

Hakim Weatherspoon

Associate Professor, Dept of Computer Science

CS 5413: High Performance Computing and Networking

March 10, 2017

Network Queueing

TCP congestion control and avoidance attempts to share bandwidth and prevent buffering (queueing)

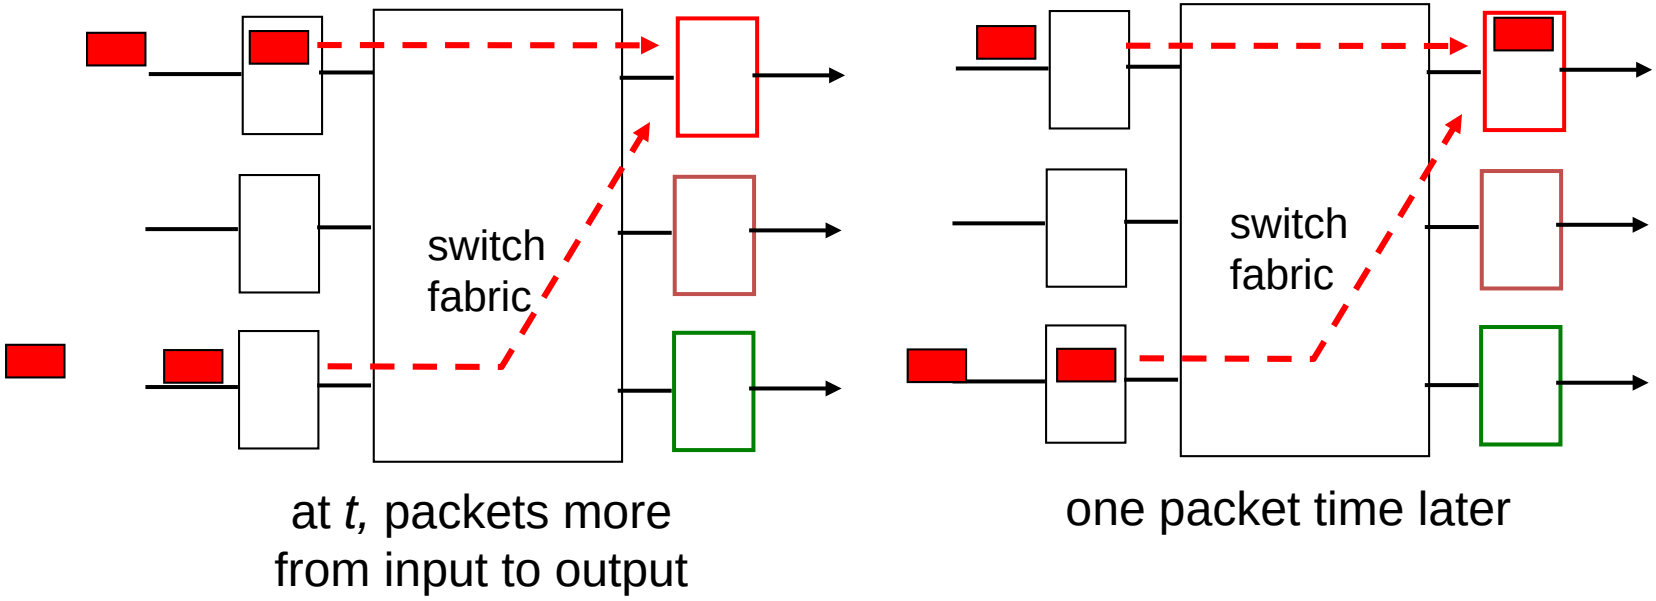

- Ideally, each flow sends at its "fair share", 1/n the network path capacity
- TCP uses packet drops to signal congestion and flows adjust rates, AIMD
- How do we proactively estimate available bandwidth?

Network Queueing

Available bandwidth estimation:

What instantaneous probe rate causes buffering?

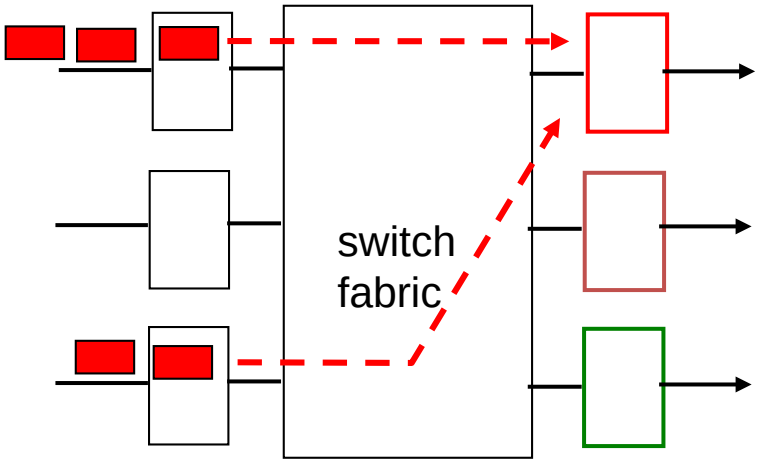

at *t*, packets more from input to output

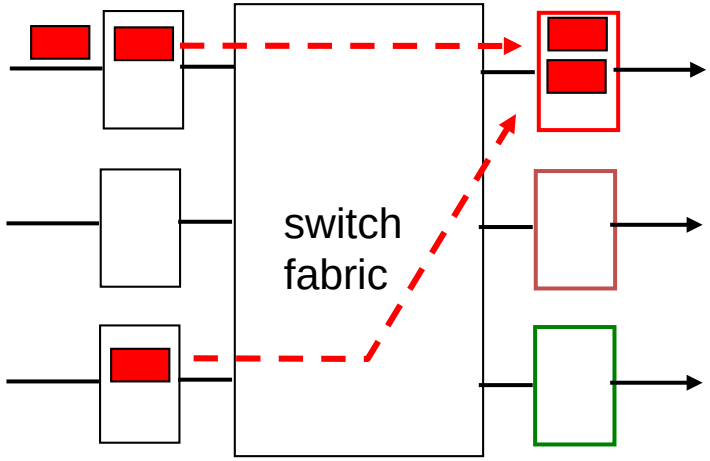

one packet time later

Active Measurement

- Narrow link: least capacity
- Tight link: least available bandwidth

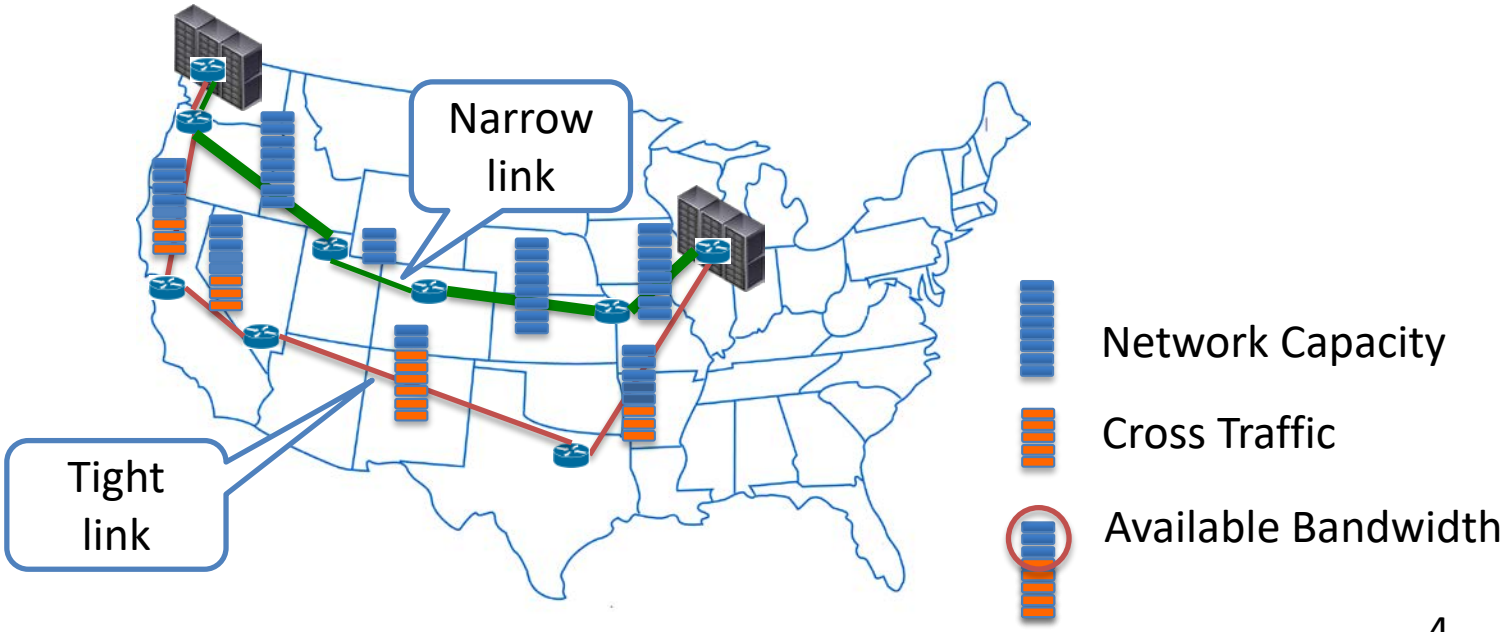

Covert timing channel

- Can we create a covert channel by modulating packet timings?
- How effective would it be? E.g. bit error rate?

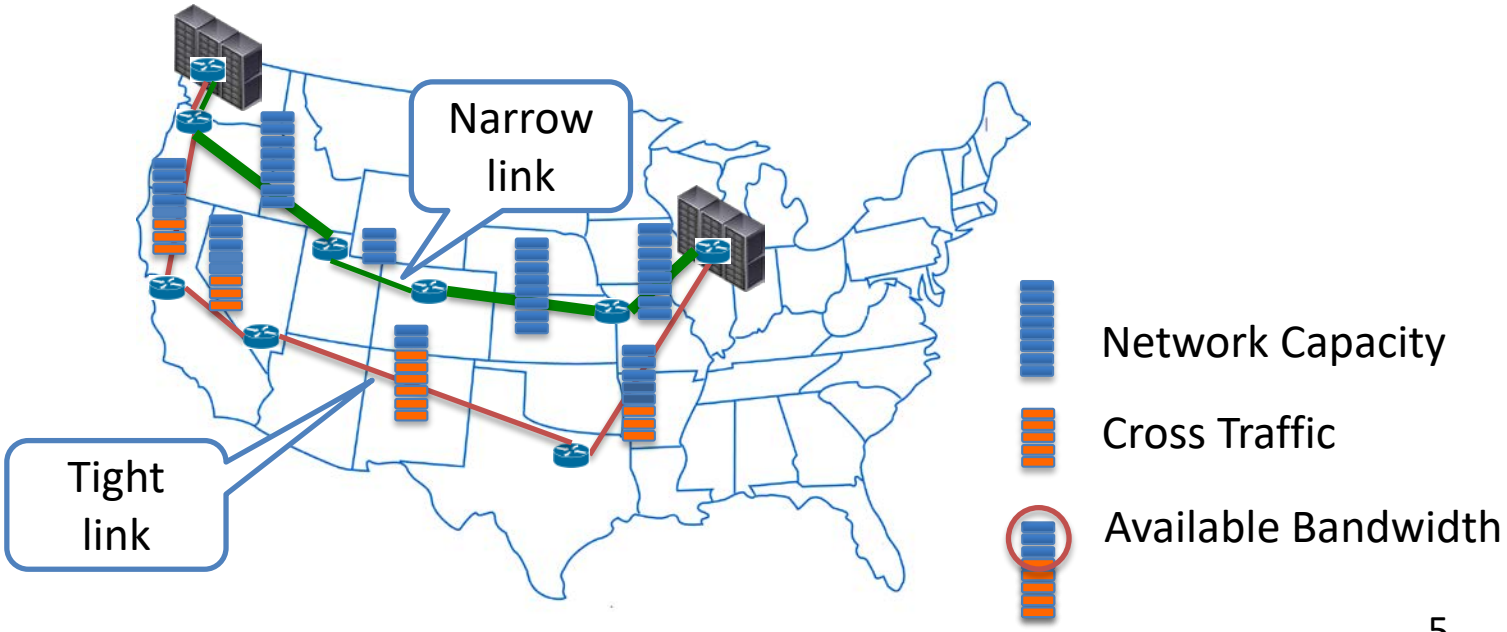

Covert Channels

- Hiding information
  - Through communication not intended for data transfer

Network Covert Channels

- Hiding information
  - Through communication not intended for data transfer
  - Using legitimate packets (Overt channel)
    - Storage Channels: Packet headers
    - Timing Channels: Arrival times of packets

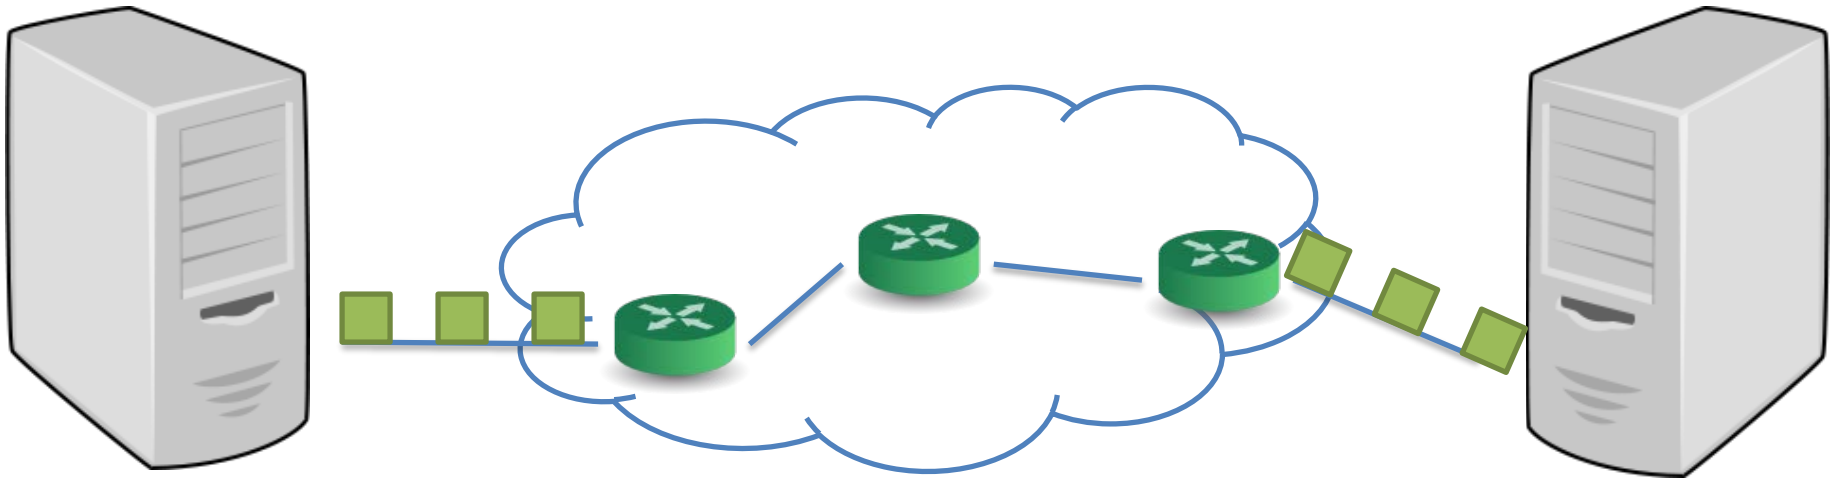

Network Covert Channels

- Hiding information
  - Through communication not intended for data transfer
  - Using legitimate packets (Overt channel)
    - Storage Channels: Packet headers
    - Timing Channels: Arrival times of packets

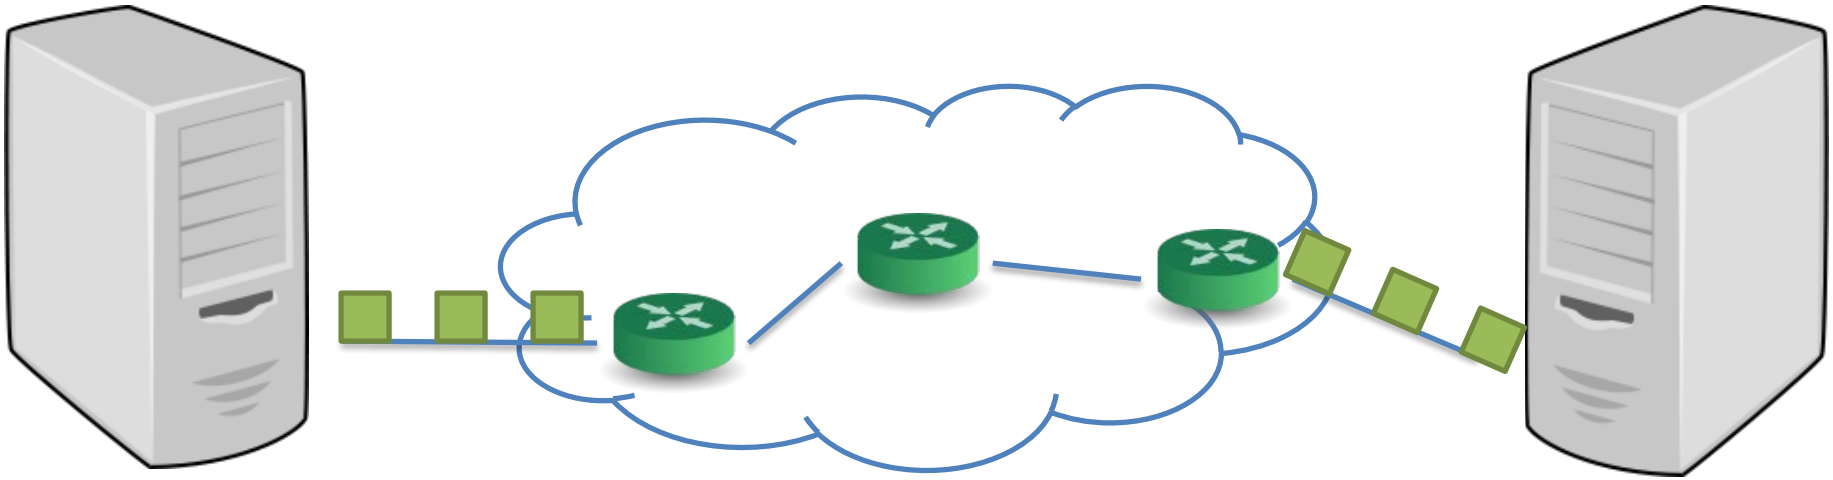

Goals of Covert Channels

- Bandwidth
  - How much information can be delivered in a second
- Robustness
  - How much information can be delivered without loss / error
- Undetectability
  - How well communication is hidden

Goals of Covert Channels

- Bandwidth
  - How much information can be delivered in a second
  - 10~100s bits per second
- Robustness
  - How much information can be delivered without loss / error
  - Cabuk'04, Shah'06
- Undetectability
  - How well communication is hidden
  - Liu'09, Liu'10

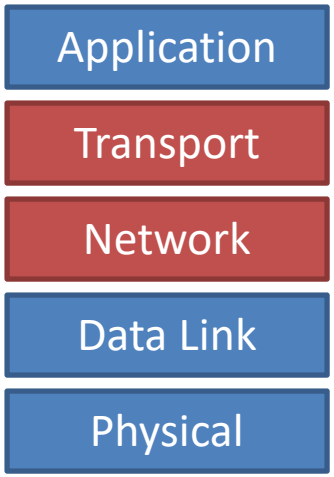

Current network covert channels
are implemented in L3~4 (TCP/IP) layers
and are *extremely* ***slow***.

Chupja: PHY Covert Channel

- Bandwidth
  - How much information can be delivered in a second
  - ~~10~100s bits per second~~ -> 10s~100s ***Kilo*** bits per second
- Robustness
  - How much information can be delivered without loss / error
  - Bit Error Rate < 10%
- Undetectability
  - How well communication is hidden
  - Invisible to detection software

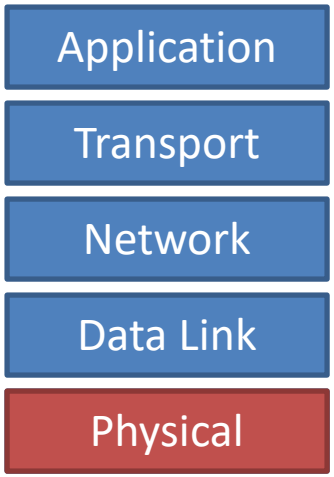

Chupja is a network covert channel
which is faster *than prior art*.

It is implemented in L1 (PHY),
robust and virtually invisible to software.

Outline

- Introduction
- Design
- Evaluation
- Conclusion

Outline

- Introduction
- Design
  - Threat Model
  - 10 Gigabit Ethernet
- Evaluation
- Conclusion

Threat Model

Passive Adversary

Commodity Server
Commodity NIC

Application
Transport
Network
Data Link
Physical

Sender

Application
Transport
Network
Data Link
Physical

Receiver

10 Gigabit Ethernet

- Idle Characters (/I/)

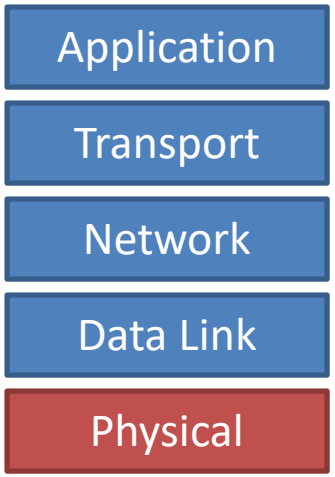

  - Each bit is ~100 picosecond wide
  - 7~8 bit special character in the physical layer
  - 700~800 picoseconds to transmit
  - Only in PHY

Terminology

- Interpacket delays (D) and gaps (G)

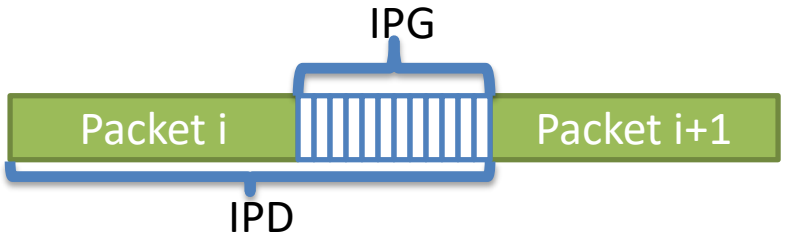

- Homogeneous packet stream

  - Same packet size,
  - Same IPD (IPG),
  - Same destination

Chupja: Design

- Homogeneous stream

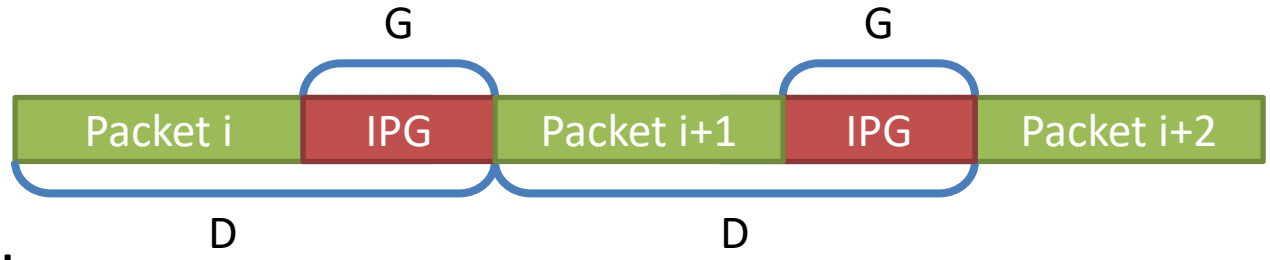

- Sender

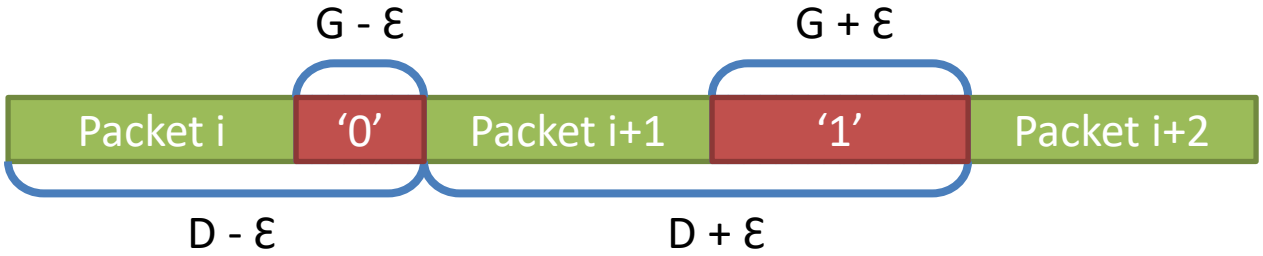

- Receiver

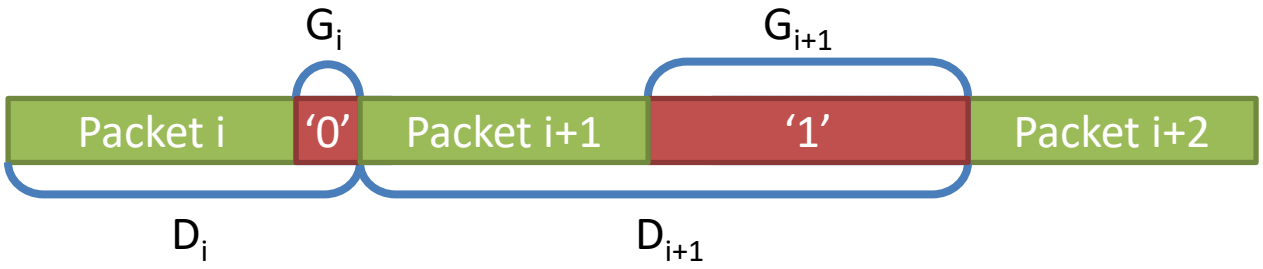

Chupja: Design

- With shared *G*
  - Encoding '1': $G_i = G + \varepsilon$
  - Encoding '0': $G_i = G - \varepsilon$

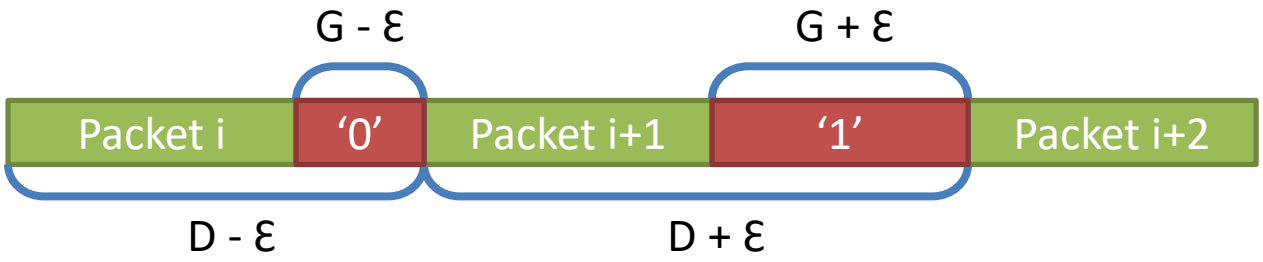

Implementation

- SoNIC [NSDI '13]
  - Software-defined Network Interface Card
  - Allows control and access *every bit* of PHY
    - In realtime, and in software
- 50 lines of C code addition

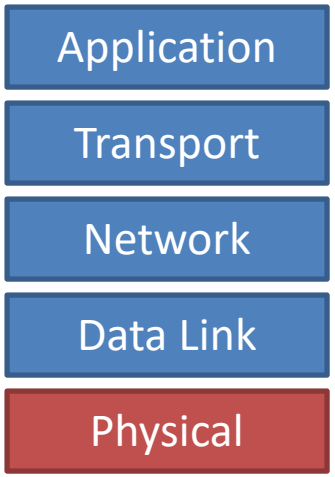

Outline

- Introduction
- Design
- Evaluation
  - Bandwidth
  - Robustness
  - Undetectability
- Conclusion

Evaluation

- What is the *bandwidth* of *Chupja*?
- How *robust* is *Chupja*?
  - *Why* is *Chupja* robust?
- How *undetectable* is *Chupja*?

What is the *bandwidth* of *Chupja*?

Evaluation: Bandwidth

- Covert bandwidth equals to ***packet rate*** of overt channel

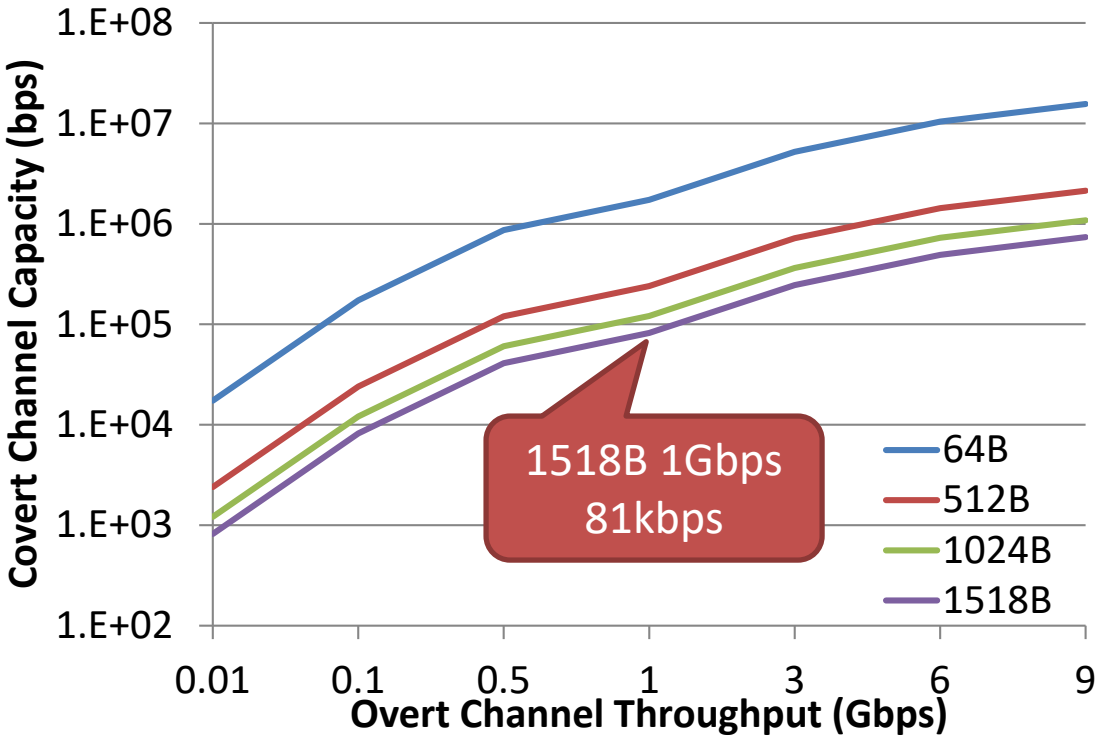

How *robust* is *Chupja*?

Evaluation Setup

- Small Network
  - Six commercial switches
  - Average RTT: 0.154 ms
- National Lambda Rail
  - Nine routing hops
  - Average RTT: 67.6ms
  - 1~2 Gbps External Traffic

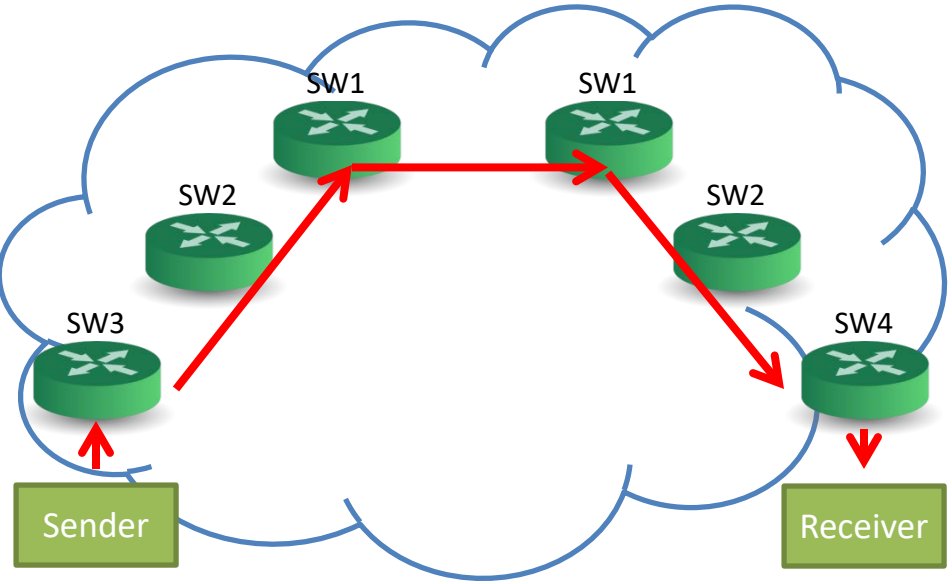

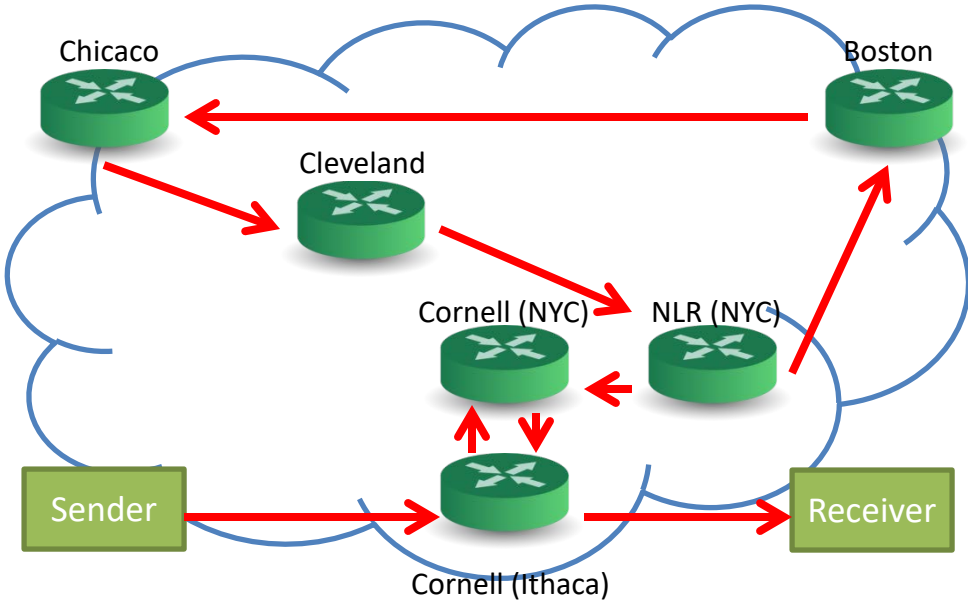

Evaluation: Robustness

- Overt Channel at 1 Gbps (D = 12211ns, G=13738 /I/s)
- Covert Channel at 81 kbps

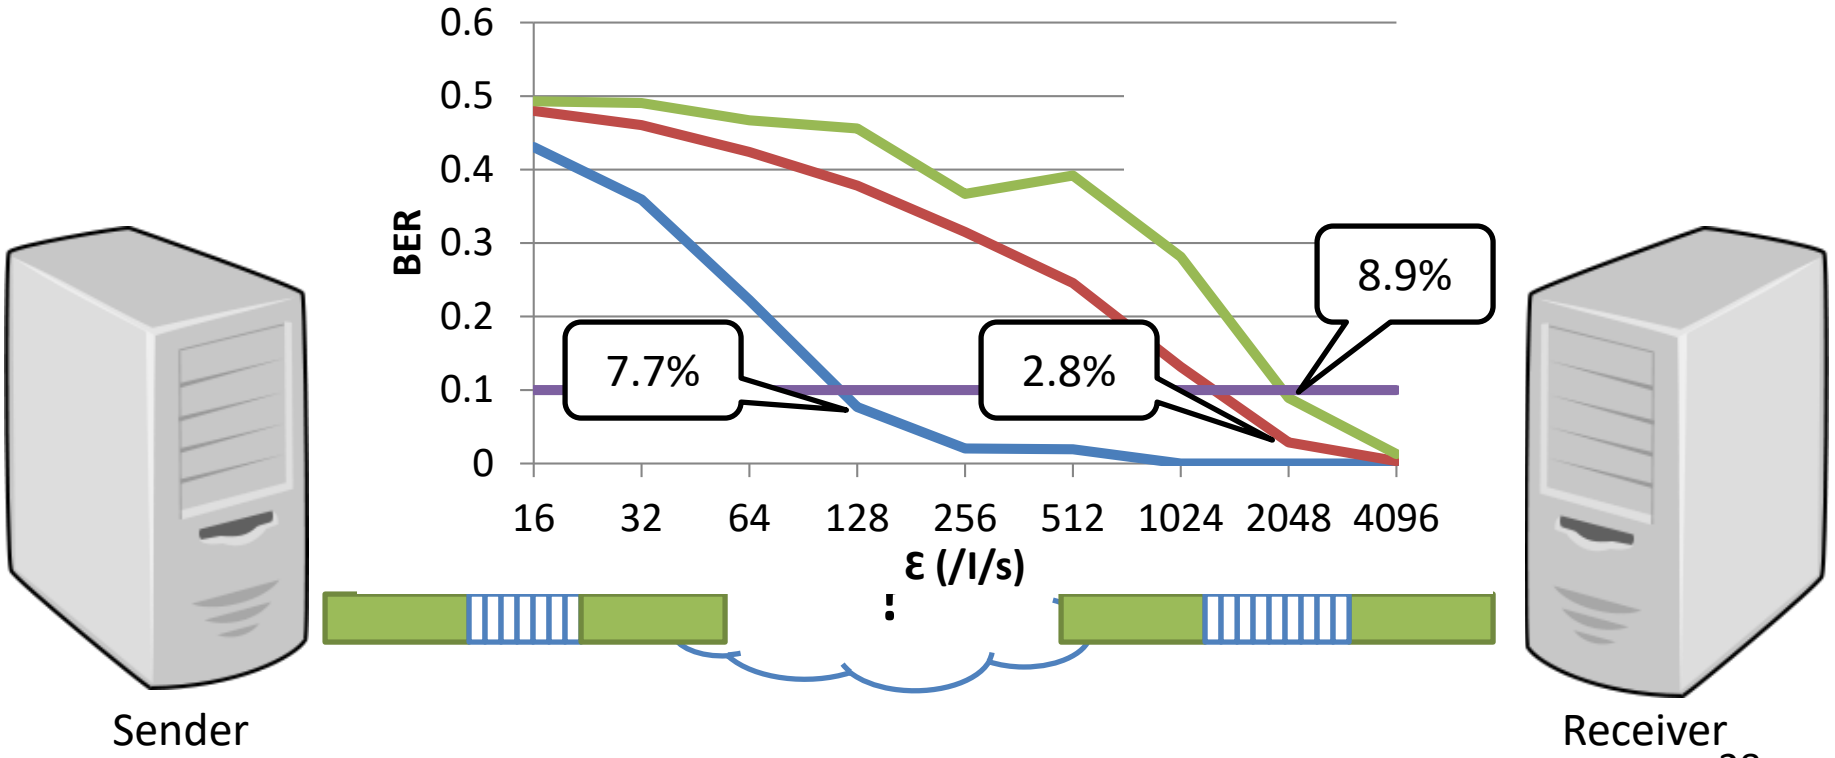

Evaluation: Robustness

- Overt Channel at 1 Gbps (D = 12211ns, G=13738 /I/s)
- Covert Channel at 81 kbps
- *Modulating IPGS at 1.6us scale (=2048 /I/s)*

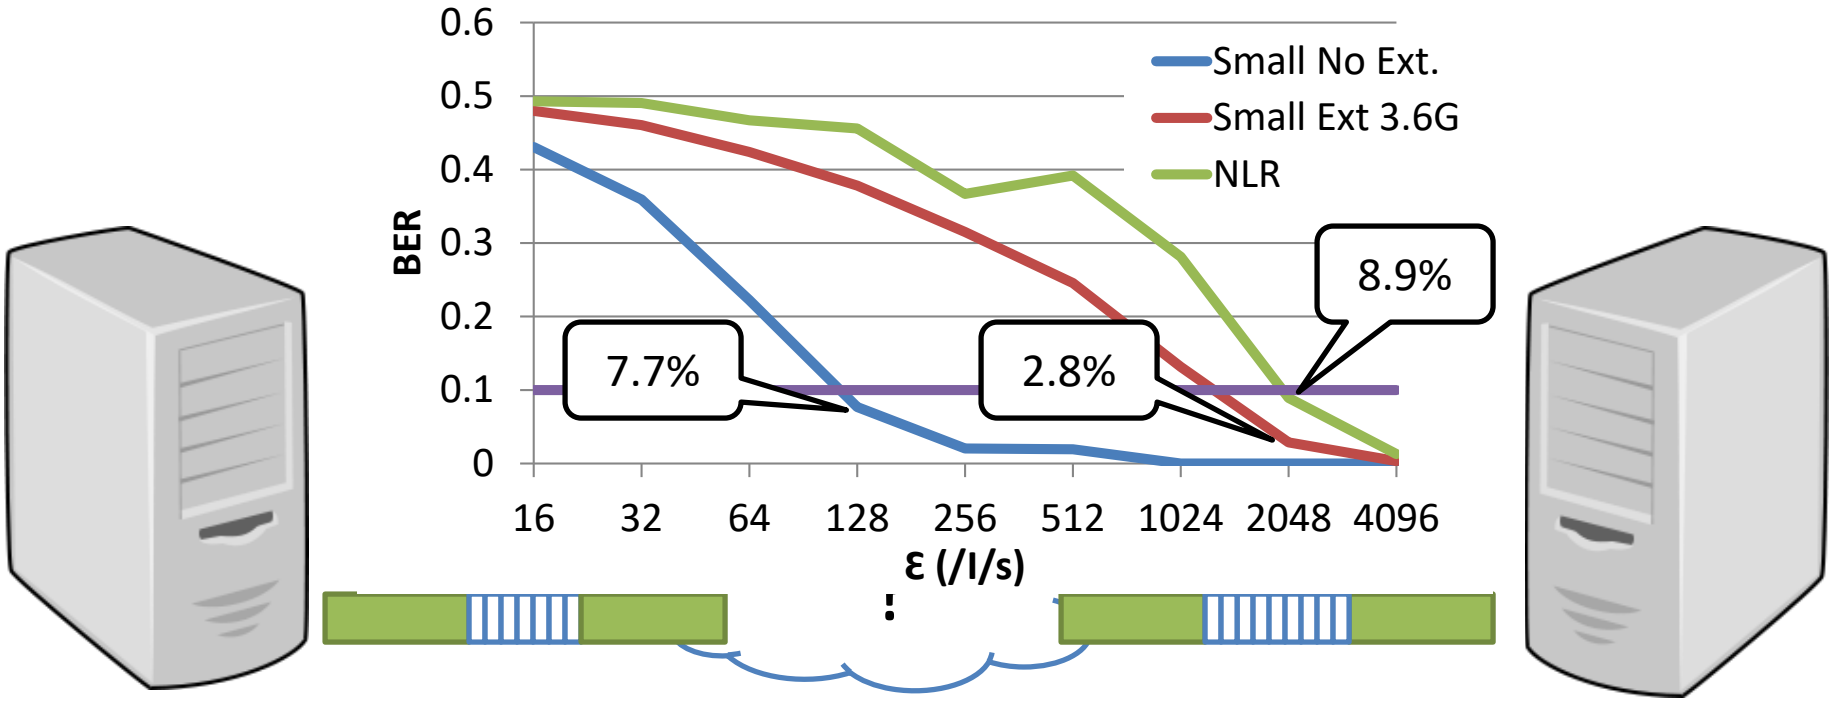

Why is *Chupja* robust?

Evaluation: Why?

- Switches do not add significant perturbations to IPDs
- Switches treat ‘1’s and ‘0’s as *uncorrelated*
  - *Over multiple hops* when there is *no external traffic*.
  - *With external traffic*

Evaluation: Why?

- Switches do not add significant perturbations to IPDs
- Switches treat '1's and '0's as *uncorrelated*
  - *Over multiple hops* when there is *no external traffic*.
  - *With external traffic*

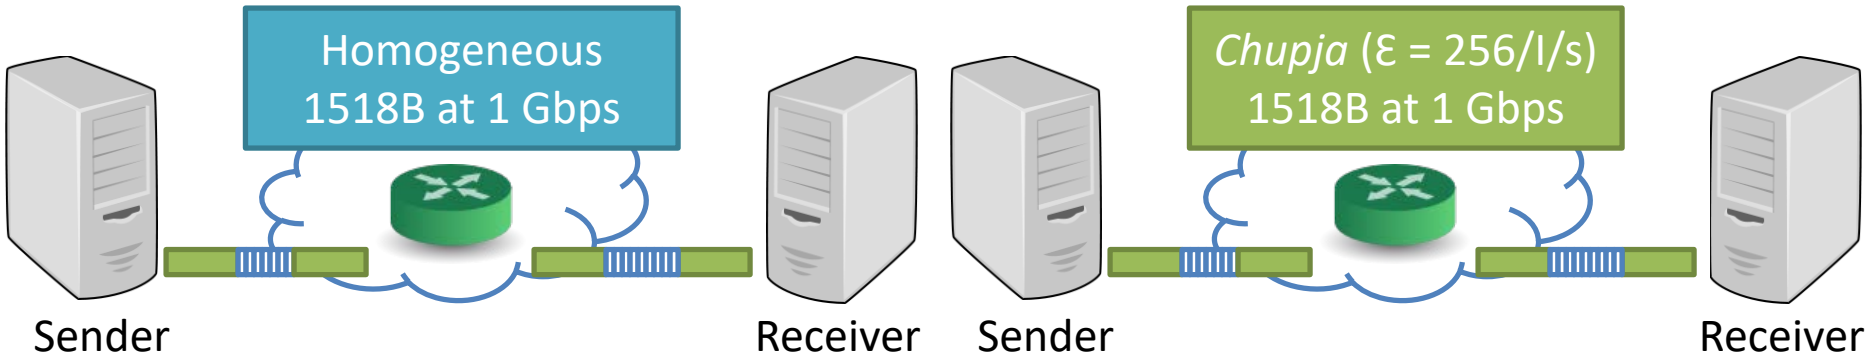

Evaluation: Why?

- Switches do not add significant perturbations to IPDs
- Switches treat encoded ‘0’ and ‘1’ as uncorrelated
  - *Over multiple hops* when there is *no* external traffic.

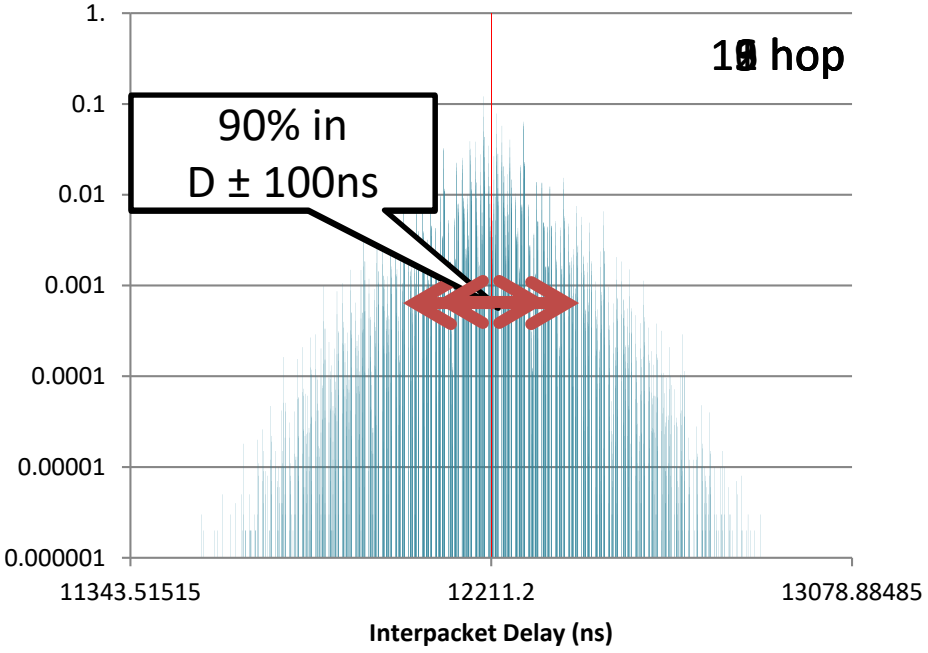

Homogeneous stream

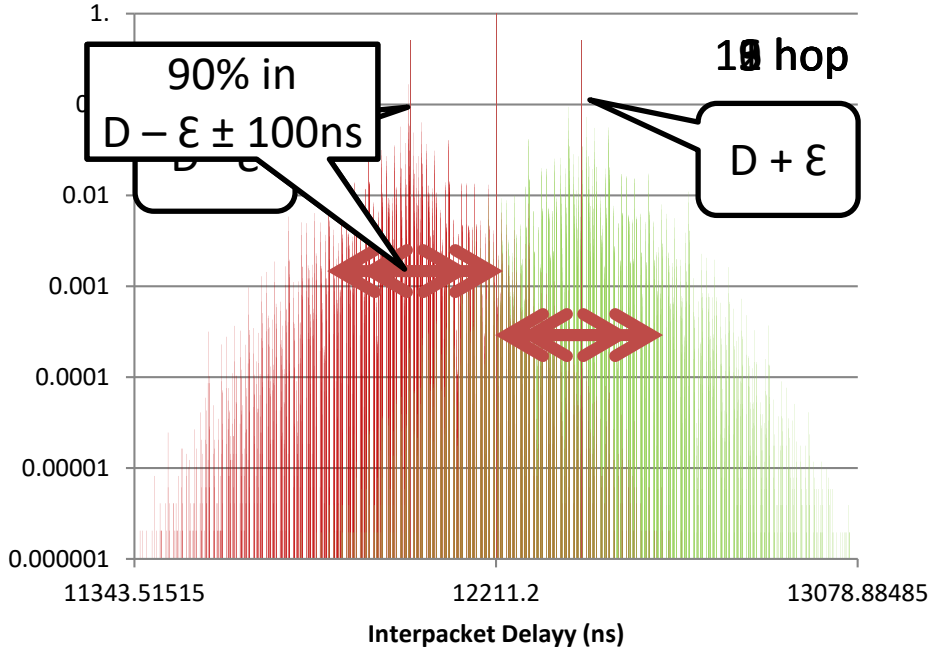

Chupja stream (Ɛ=256/l/s)

Evaluation: Why?

- Most of IPDs are within some range from original IPD
  - Even when there is *external traffic*.

Ɛ (/I/s) (ns)	256 (=204.8ns)	512 (=409.6)	1024 (=819.2)	2048 (=1638.4)	4096 (=3276.8)
BER					

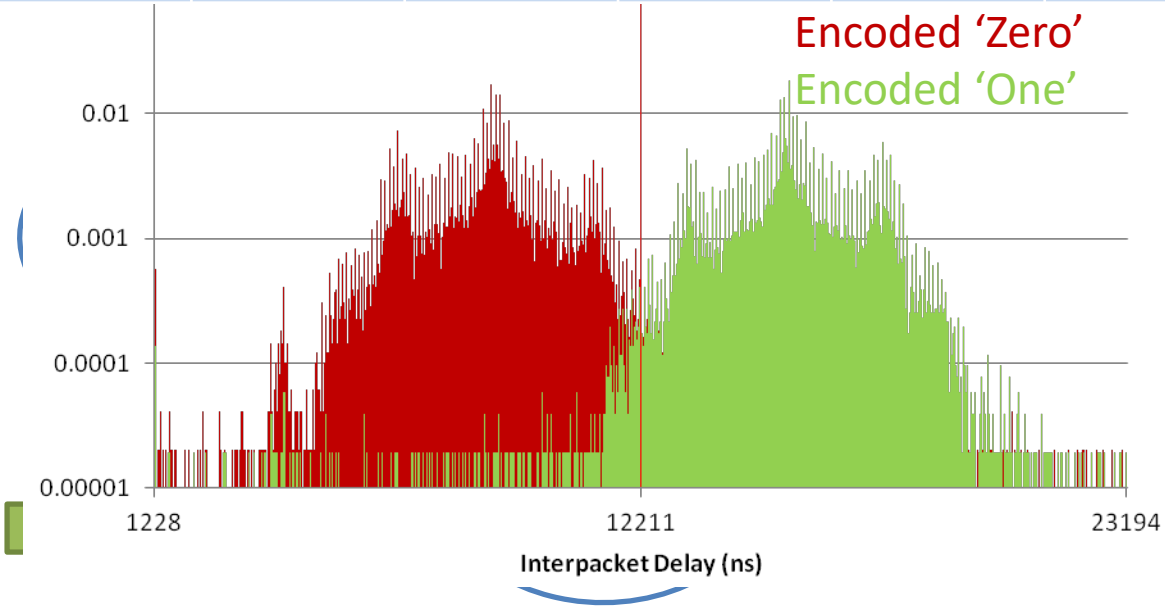

Sender

Receiver

Evaluation: Why?

- Switches do not add significant perturbations to IPDs
- Switches treat '1's and '0's as *uncorrelated*
  - *Over multiple hops* when there is *no external traffic*.
  - *With external traffic*

With sufficiently large Ɛ,
the interpacket spacing holds throughout
the network, and BER is less than 10%

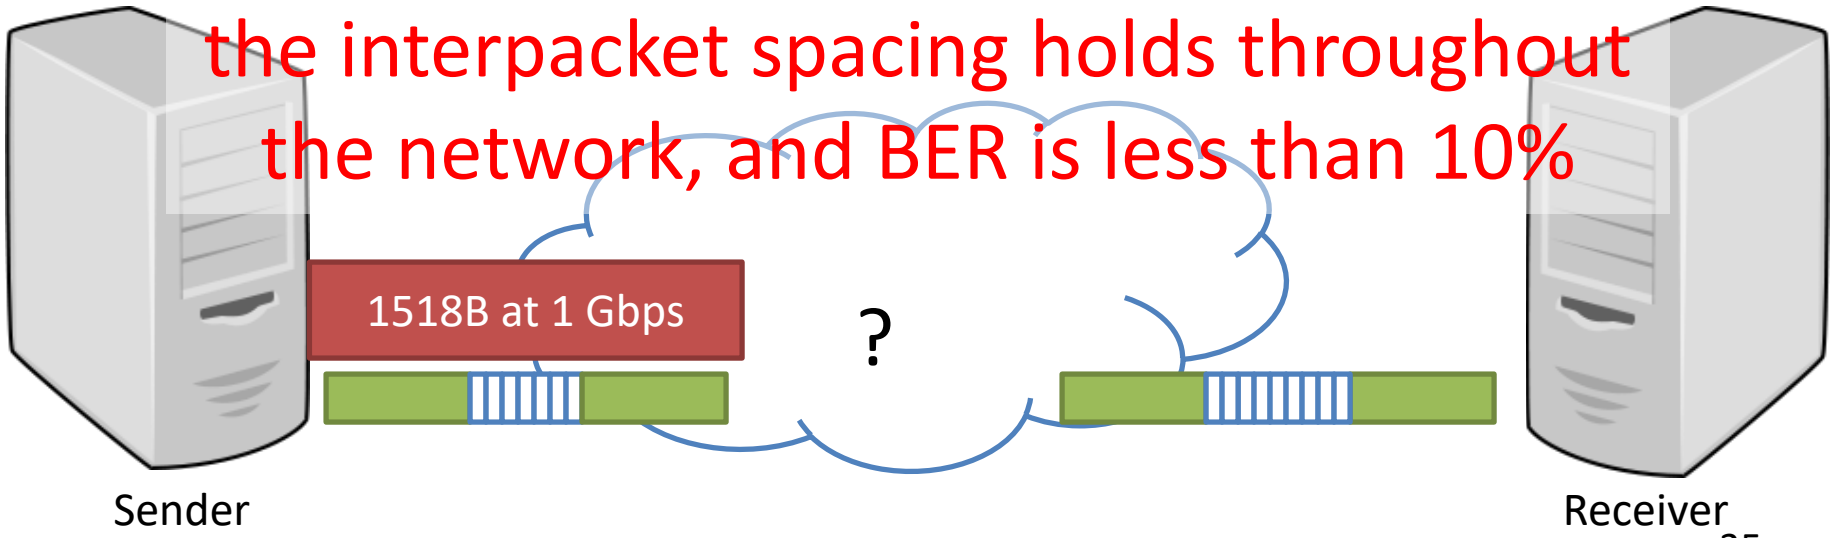

How *undetectable* is *Chupja*?

Evaluation: Detection Setup

- Commodity server with 10G NIC
  - Kernel timestamping

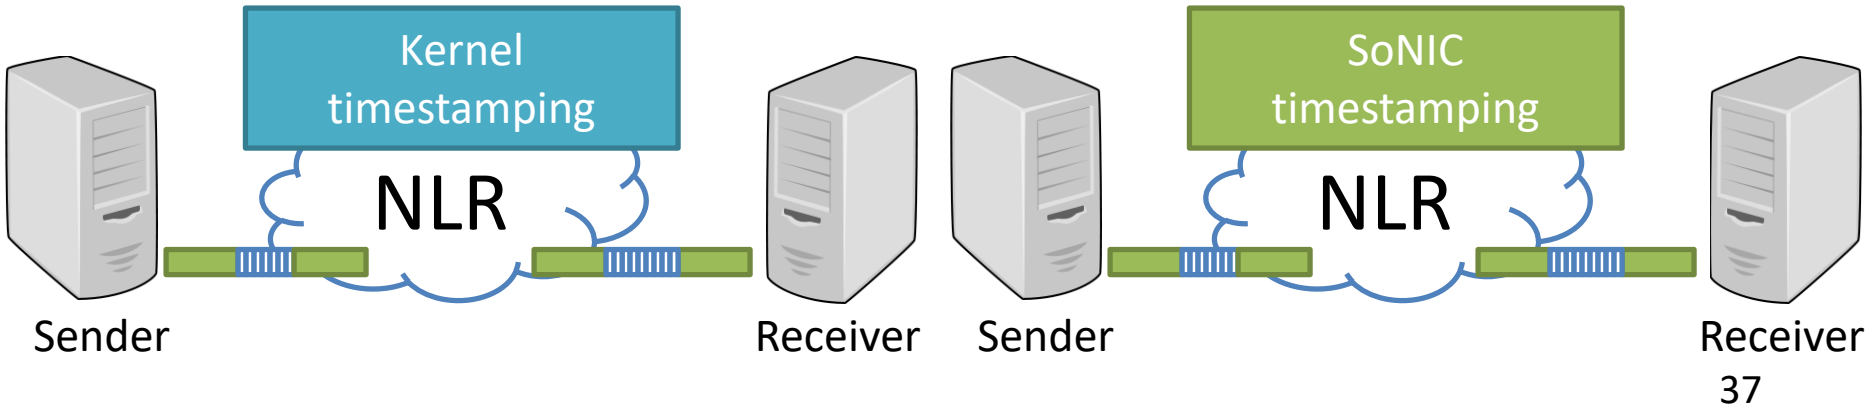

Evaluation: Detection

- *Adversary cannot detect patterns of Chupja*

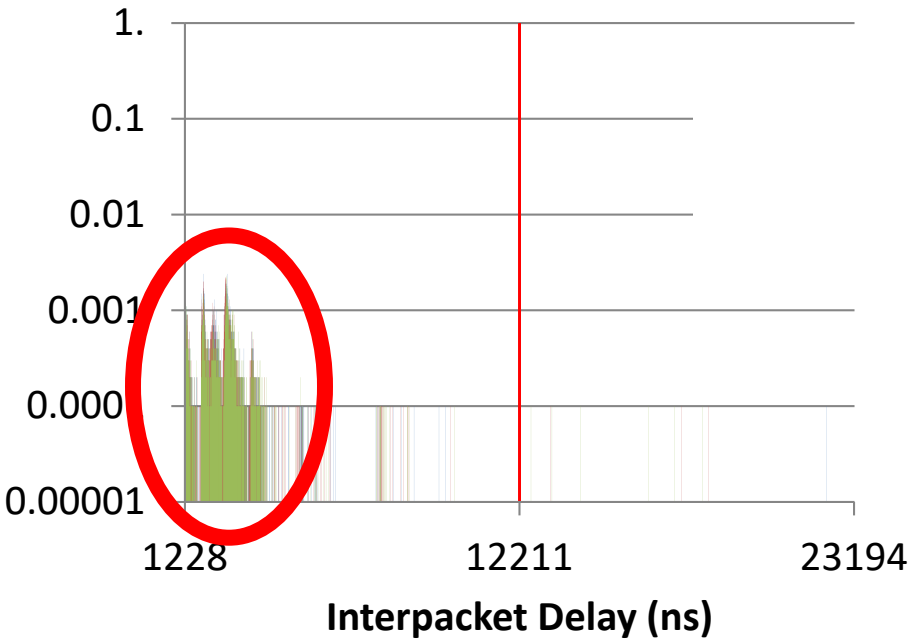

Kernel Timestamping

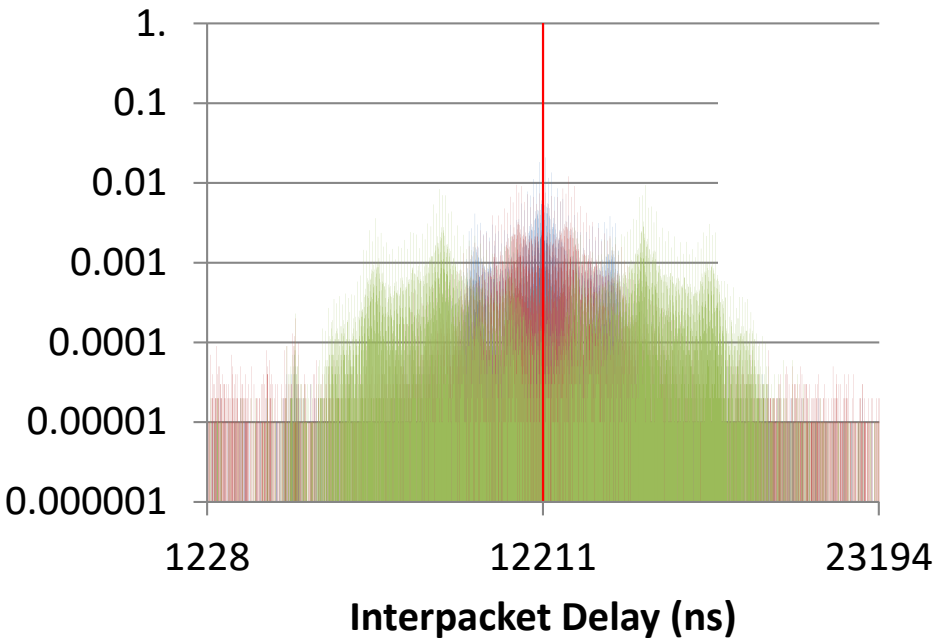

SoNIC Timestamping

Evaluation: Summary

- What is the *bandwidth* of *Chupja*?
  - 10s~100s Kilo bits per second
- How *robust* is *Chupja*?
  - BER < 10% over NLR
  - *Why* is *Chupja* robust?
    - Sufficiently large Ɛ holds throughout the network
- How *undetectable* is *Chupja*?
  - Invisible to software

Before Next time

- Project
  - Continue to make progress.
  - **Intermediate project report due Mar 22. BOOM proposal due Mar 29.**
- HW2
  - Chat Server
  - **Due this Sunday, March 12**
- Monday, bring your laptop
- Check website for updated schedule

Where are we in the semester?

- Overview and Basics
  - Overview
  - Basic Switch and Queuing (today)
  - Low-latency and congestion avoidance (DCTCP)
- Data Center Networks
  - Data Center Network Topologies
  - Software defined networking
    - Software control plane (SDN)
    - Programmable data plane (hardware [P4] and software [Netmap])
  - Rack-scale computers and networks
  - Disaggregated datacenters
  - Alternative Switching Technologies
  - Data Center Transport
  - Virtualizing Networks
  - Middleboxes
- Advanced topics