PHY Covert Channels: Can you see the Idles?

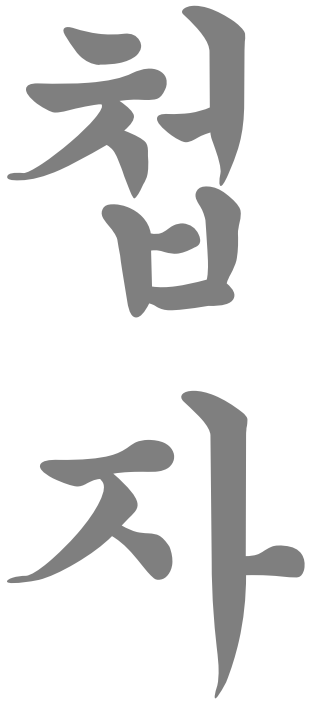

Chupja

Prof. Hakim Weatherspoon
Joint with Ki Suh Lee, Han Wang
Cornell University

Royal Holloway, University of London
June 27, 2016

첩자 (chupja)


Goguryeo 고구려
Baekje 백제
Silla 신라

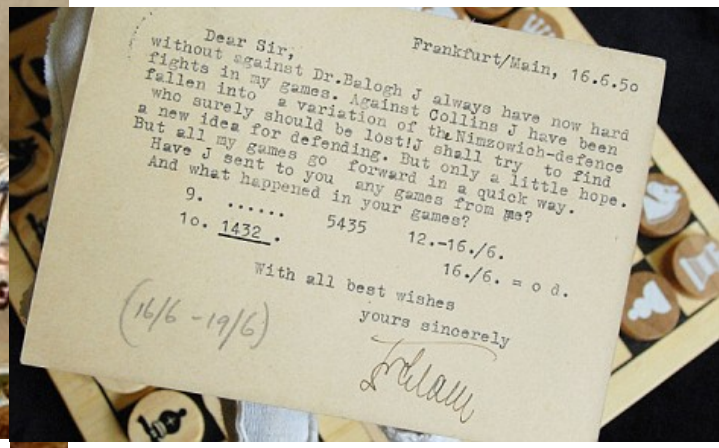
Dear Sir,
Frankfurt/Main, 16.6.5o
without against Dr.Balogh J always have now hard
fights in my games. Against Collins J have been
fallen into a variation of the Nimzowich-defence
who surely should be lost! J shall try to find
a new idea for defending. But only a little hope.
But all my games go forward in a quick way.
Have J sent to you any games from me?
And what happened in your games?
9. ..... 5435 12.-16./6.
1o. 1432. 16./6. = o d.
With all best wishes
yours sincerely
(16/6 - 19/6)

Setauket ★ Long Island ★ New York ★ 1778-1781
SPIES!

Can an underground spy ring exist and thrive within the Internet?

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

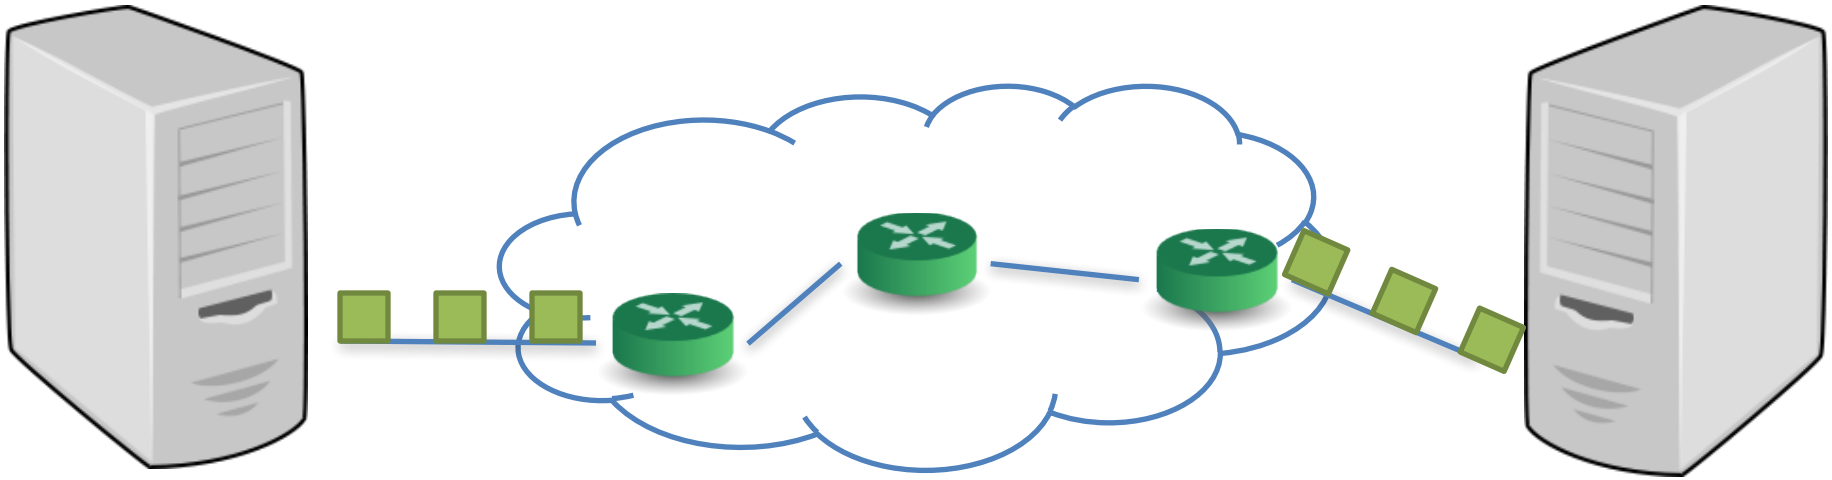

Network Covert Channels

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

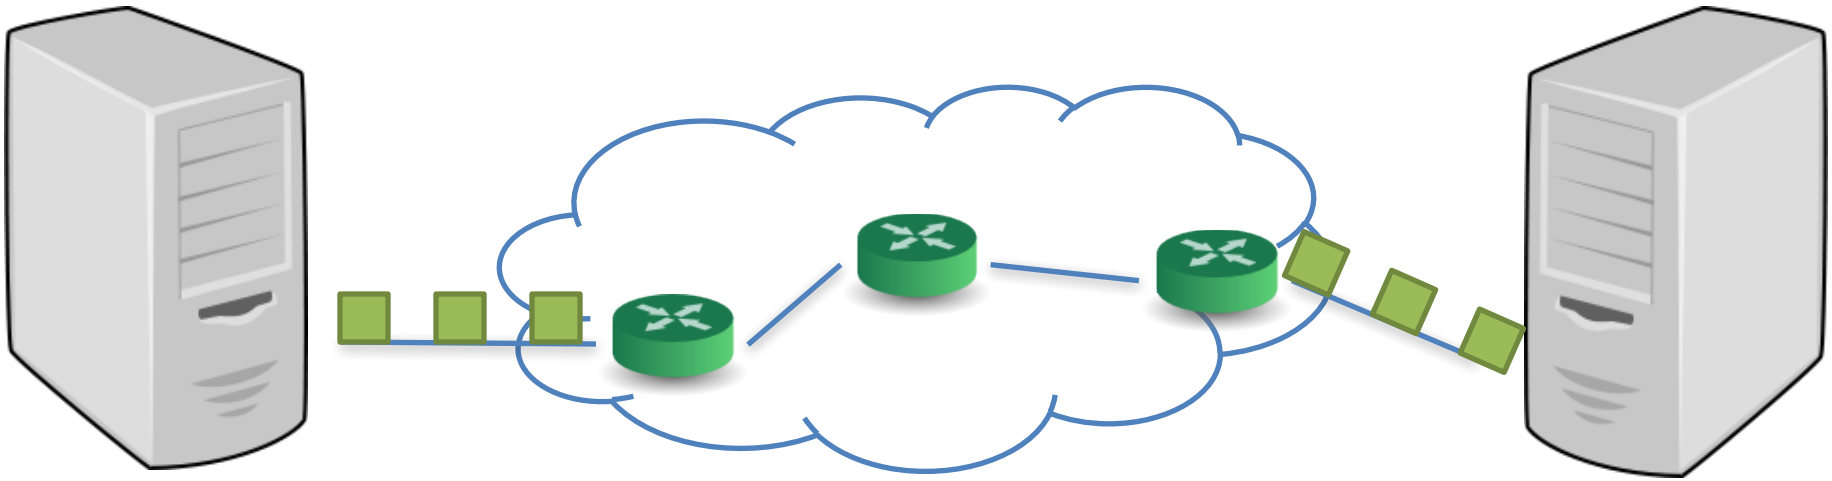

Goals of Covert Channels

- 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

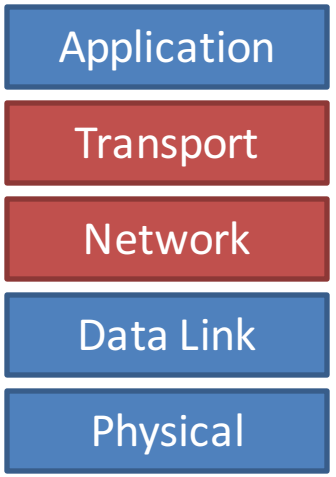

Current network covert channels
are implemented in L3~4 (TCP/IP) layers
and are *extremely* ***slow***.

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

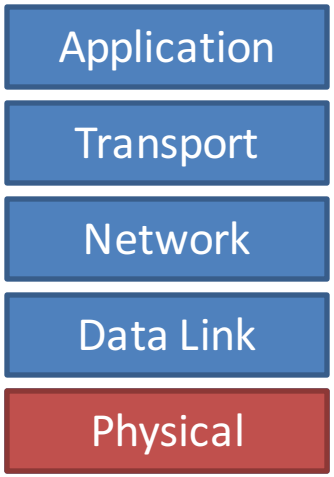

Chupja is a network covert channel
which is faster *than prior art*.

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

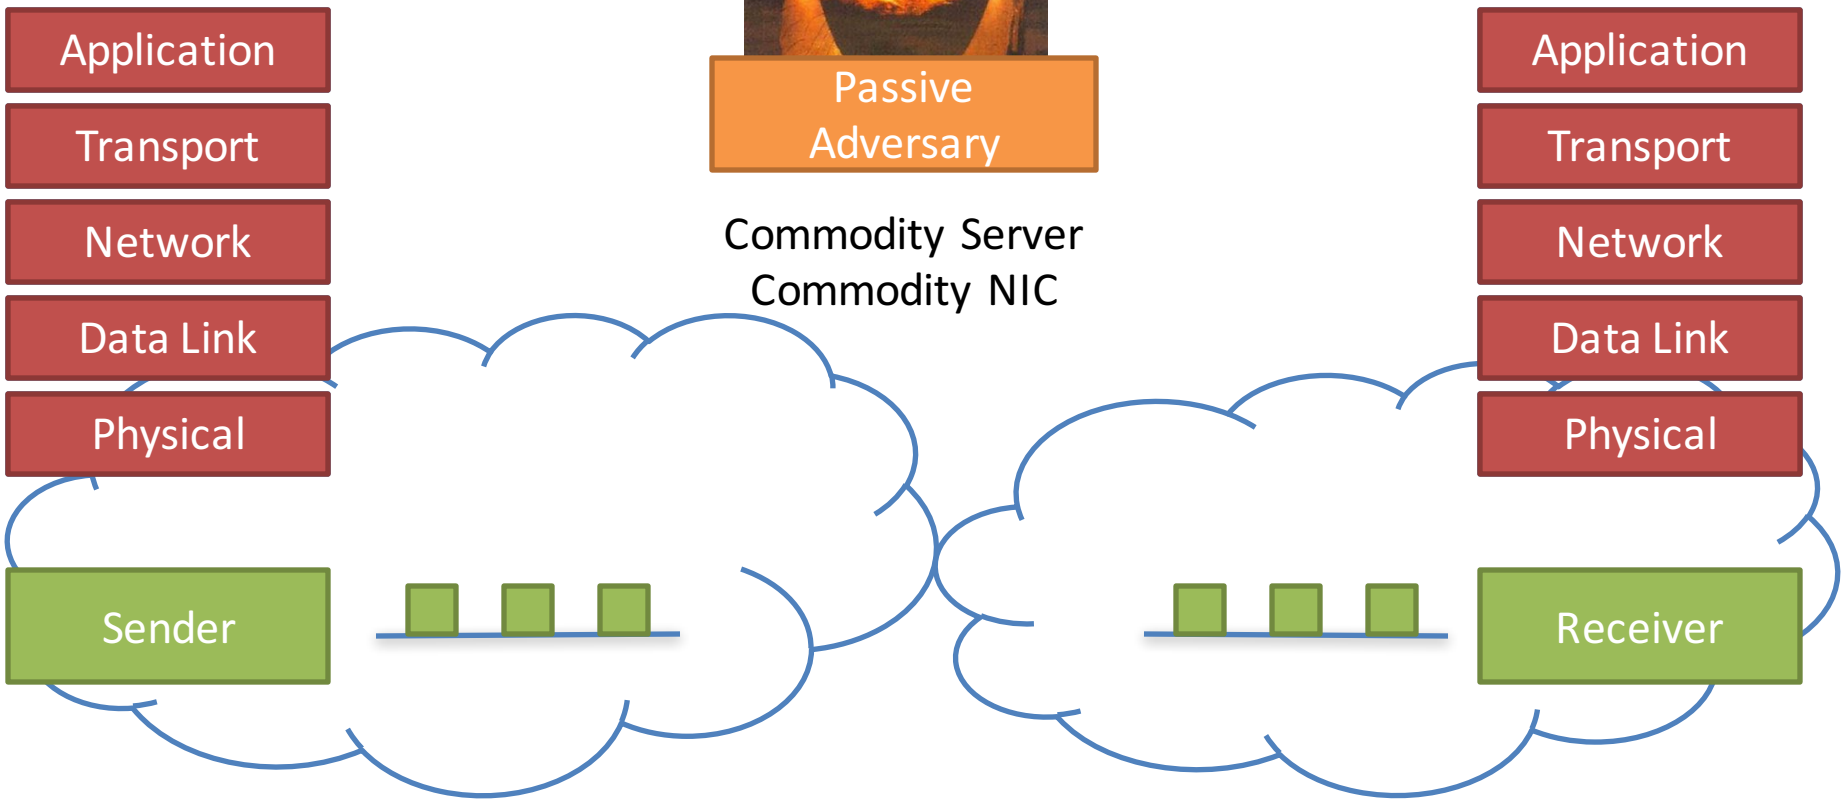

10 Gigabit Ethernet

- Idle Characters (/I/)

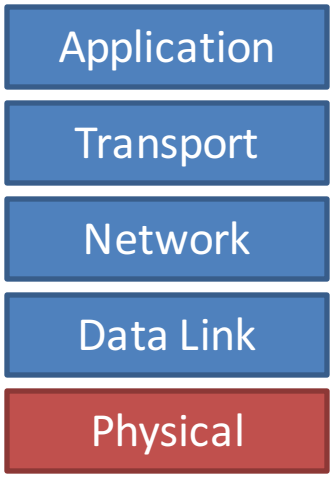

  - Each bit is ~100 picosecond wide
  - 7~8 bit special character in the physical layer
  - 700~800 picoseconds to transmit
  - Only in PHY

Terminology

- Interpacket delays (D) and gaps (G)

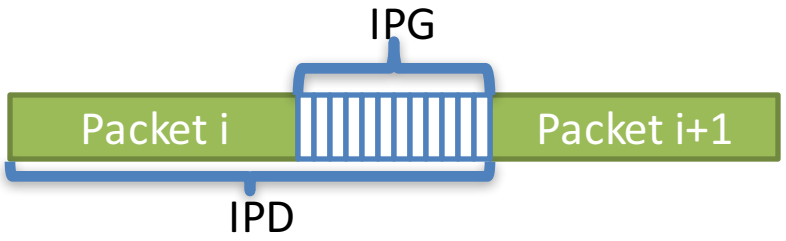

- Homogeneous packet stream

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

Chupja: Design

- Homogeneous stream

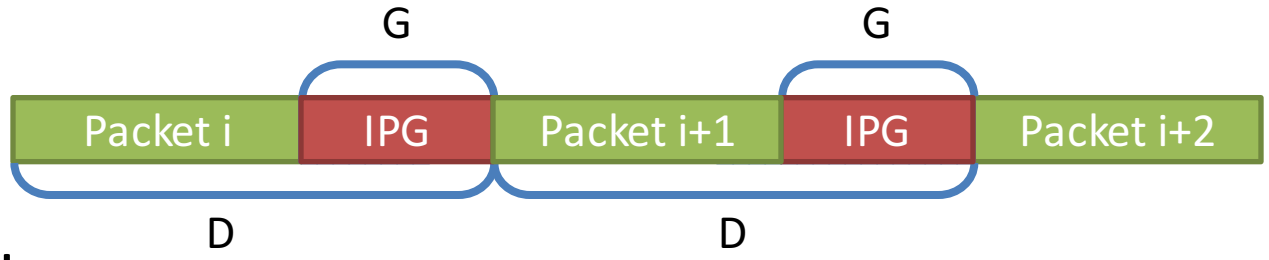

- Sender

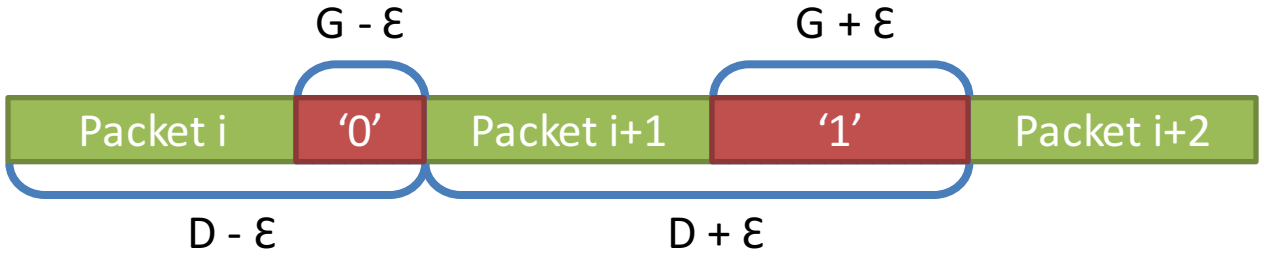

- Receiver

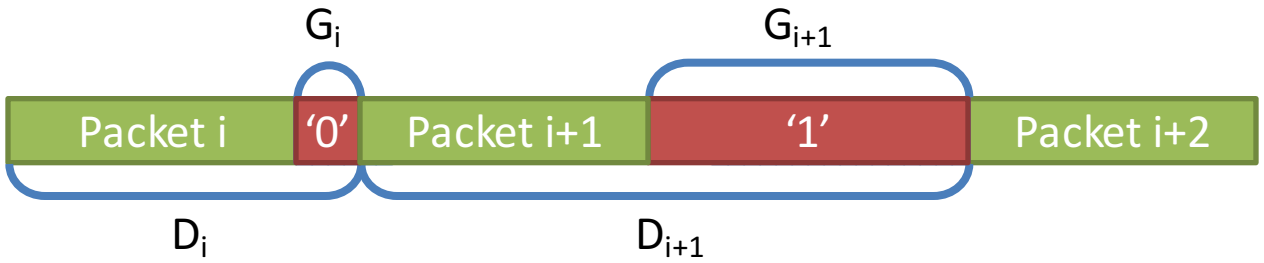

Chupja: Design

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

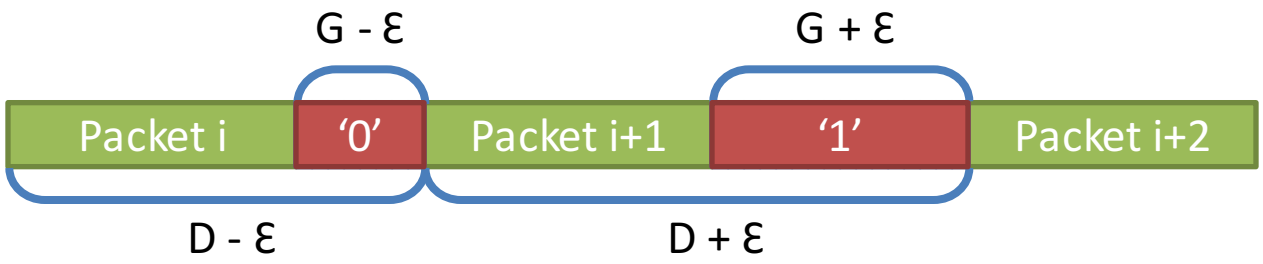

Implementation

- 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

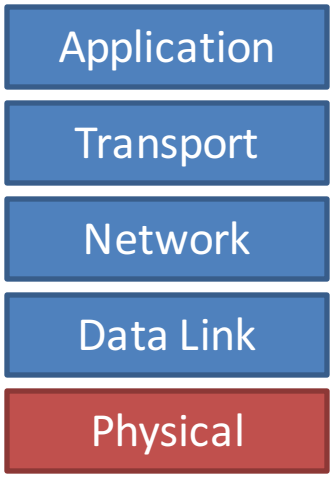

Outline

- 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

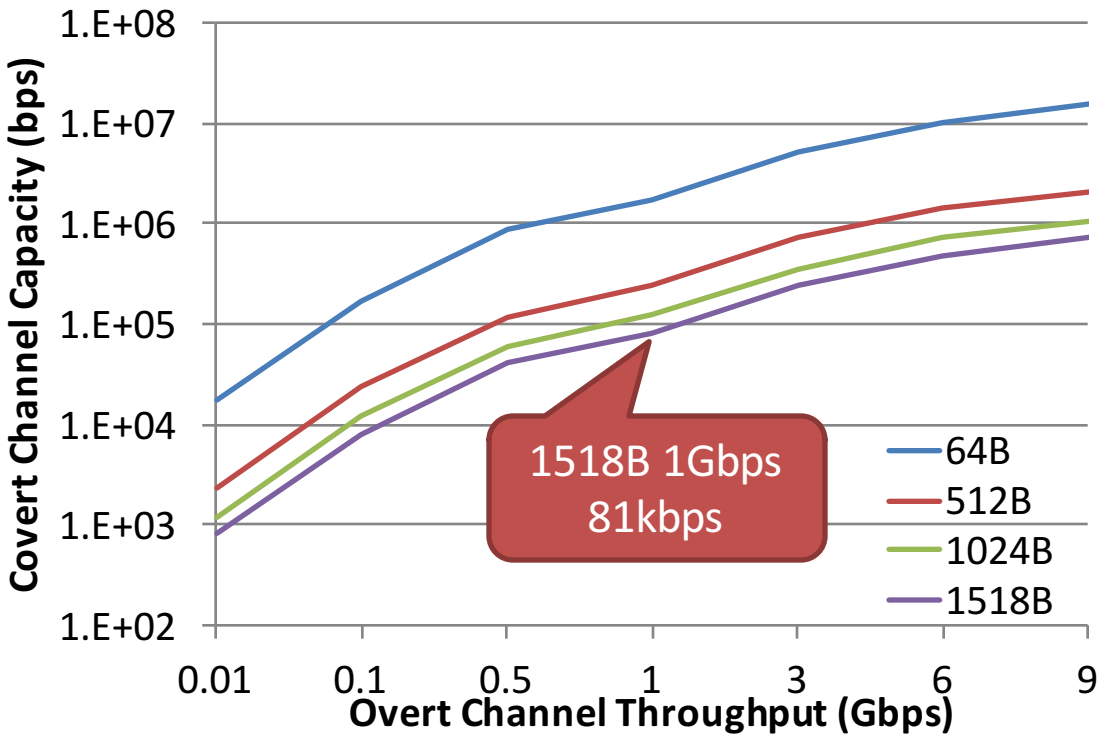

How *robust* is *Chupja*?

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

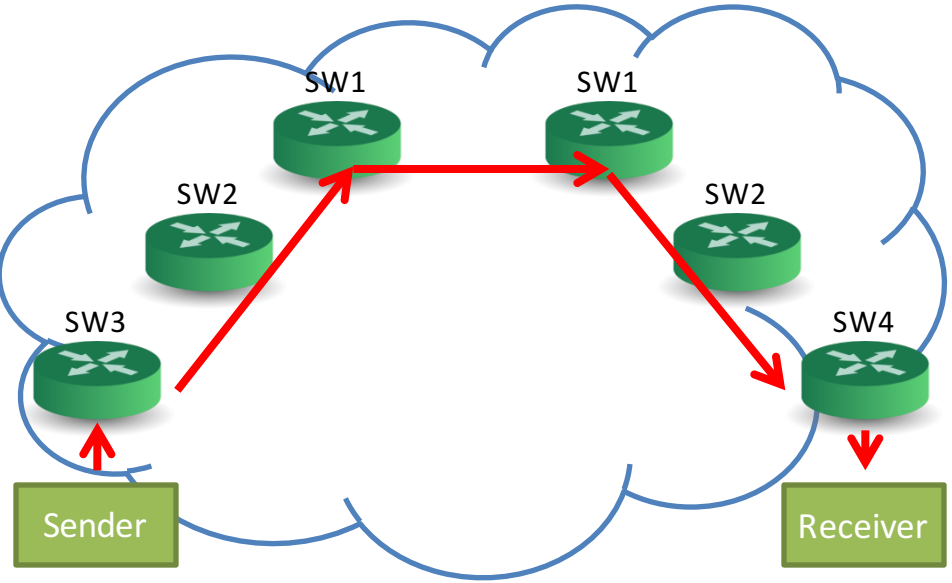

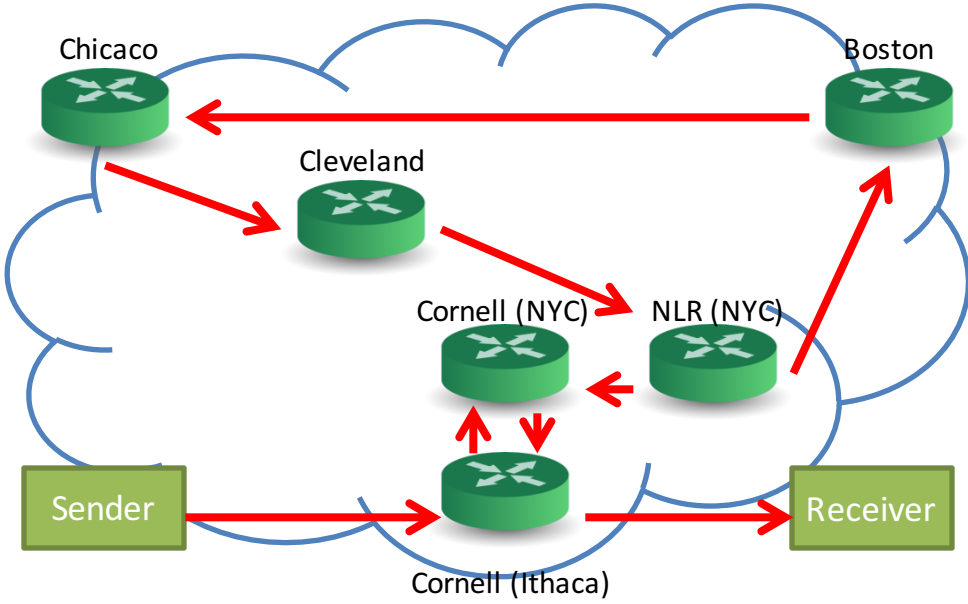

Evaluation: Robustness

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

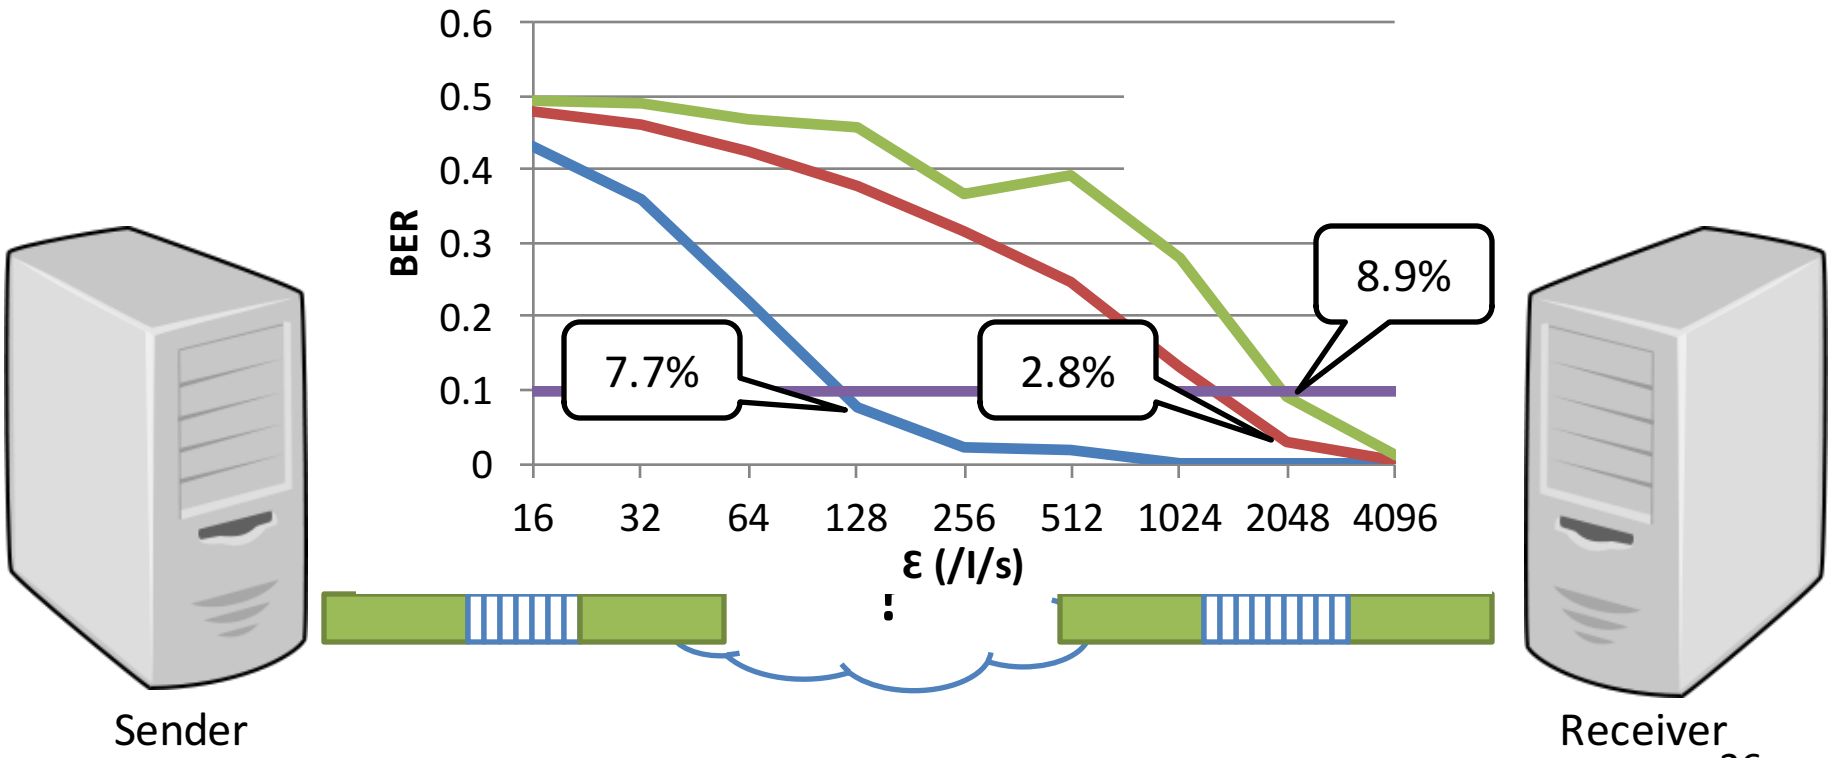

Evaluation: Robustness

- 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)*

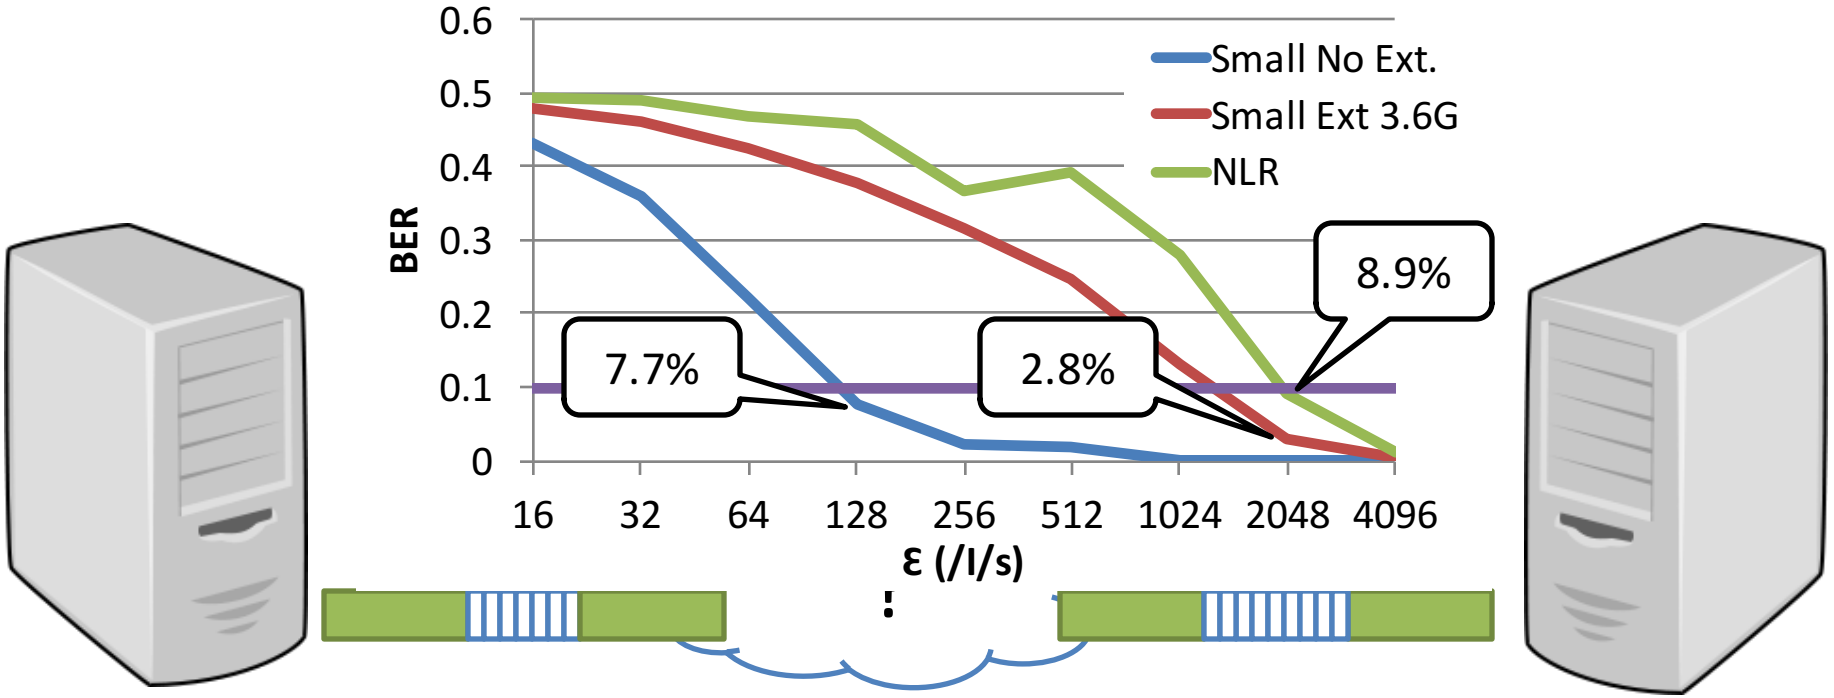

Evaluation: Why?

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

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

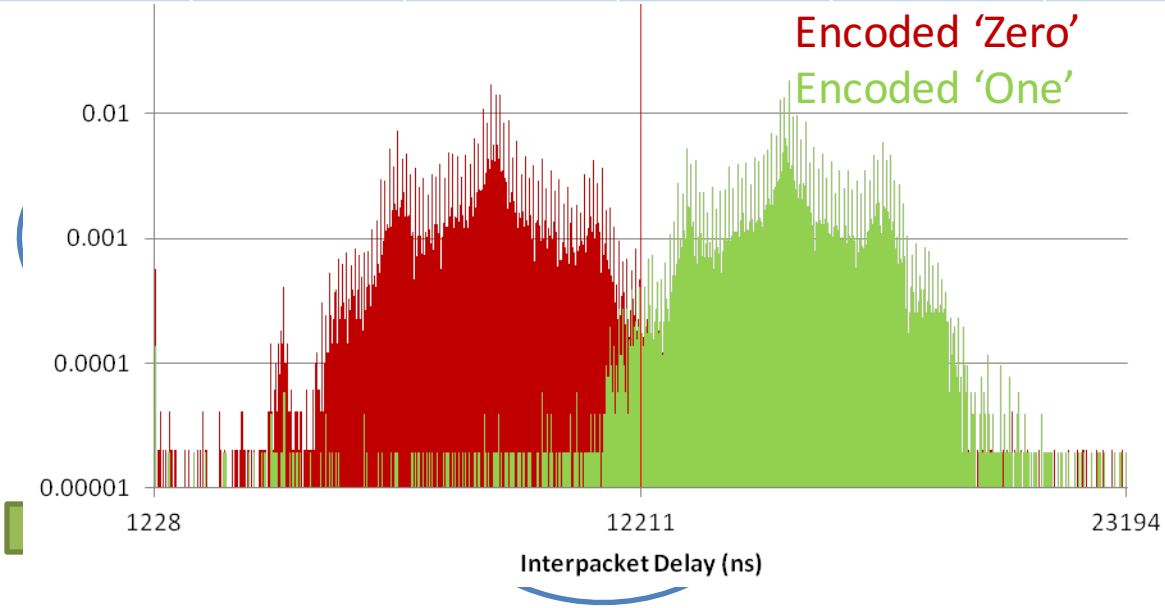

Sender

Receiver

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

Broader Context

- Why access the physical layer from software?

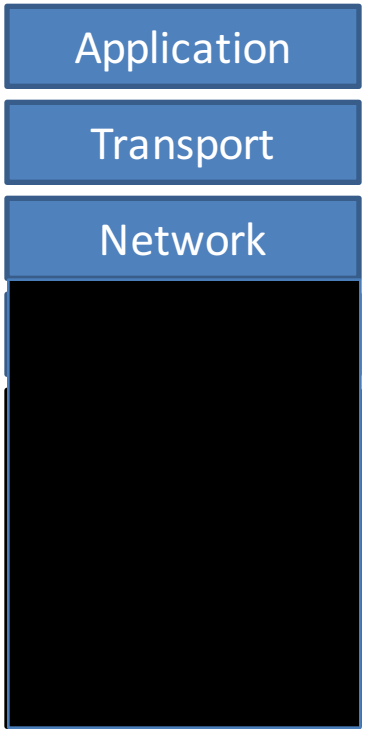

Issue:

- Programmers treat layers 1 and 2 as black box

Opportunities

- Network Measurements
- Network Monitoring/Profiling
- Network Steganography

- Can improve security, availability, and performance of the distributed systems cloud networks

Accurate available bandwidth estimation [IMC 2014]

Control at 100ps

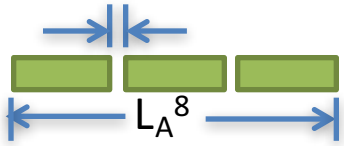

Measure at 100ps

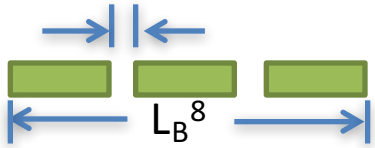

We advance the State-of-art in available bandwidth estimation because we can control and capture inter-packet spacing with exact precision.

Datacenter Time Protocol [SIGCOMM 2016]

Unprecedented, Precise, and bounded synch

- 4 clock ticks / 25 ns bounded peer-wise synchronization
- 100ns precision synchronization for an entire datacenter
- No clock differs by more than 100ns
- Free – No network traffic: Use the PHY!

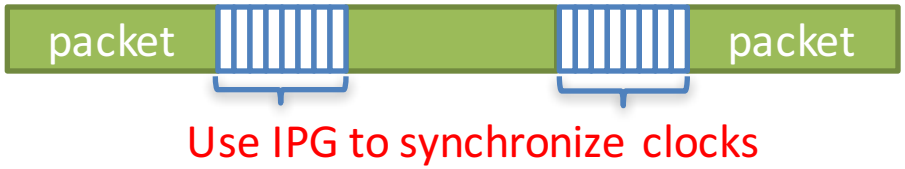

Rack-scale computing: Coordination Free Networks

- Assuming synchronized time, schedule every packet
- Every node is allocated a full time slot to a single destination
- No two nodes will be able to communicate with the same destination at the same time

	1	2	3	4
Node 1	2	3	4	5
Node 2	3	4	5	1
Node 3	4	5	1	2
Node 4	5	1	2	3
Node 5	1	2	3	4

Rack-scale computing: Coordination Free Networks

Benefits

- No network contention
- Full bisection bandwidth
  - Direct connect topology
  - Route through one random intermediate node
- Bounded latency
- Low power

	1	2	3	4
Node 1	2	3	4	5
Node 2	3	4	5	1
Node 3	4	5	1	2
Node 4	5	1	2	3
Node 5	1	2	3	4

P4FPGA [http://p4fpga.org]

- P4: Programming Protocol-Independent Packet Processors
- Use P4 to describe many different network applications
- Compile P4 to many different FPGA platforms

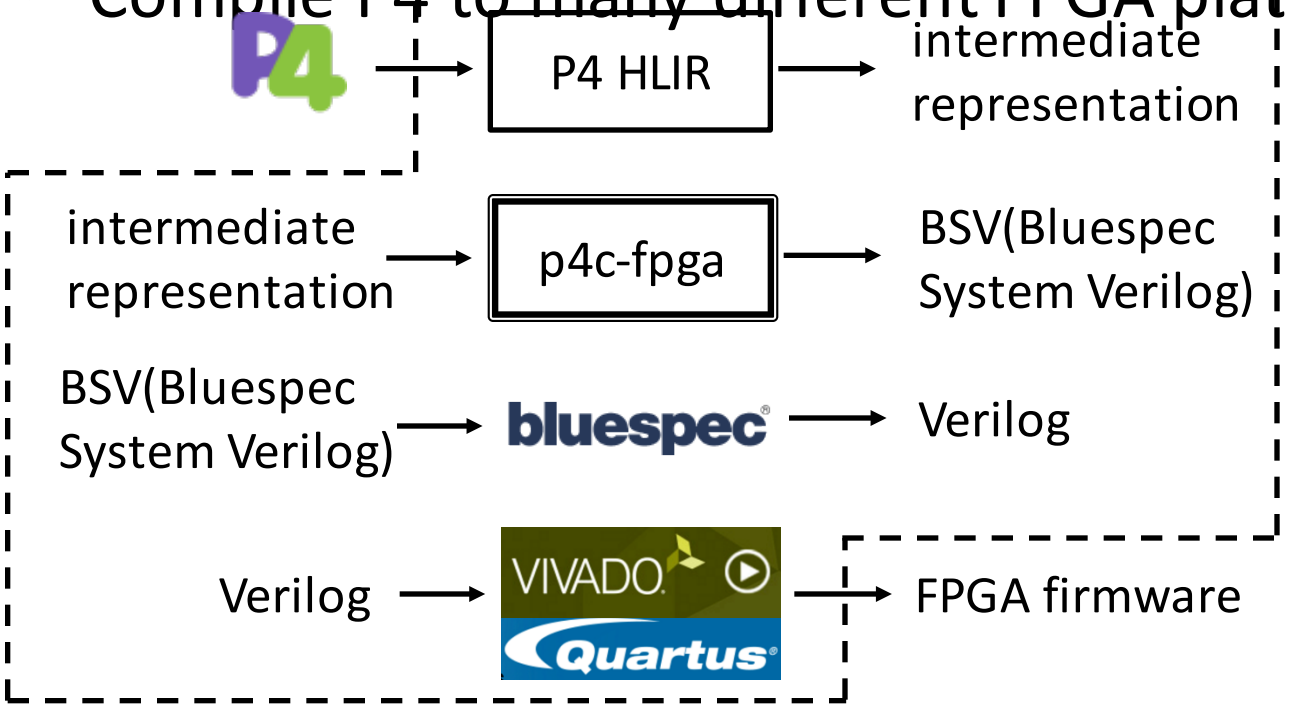

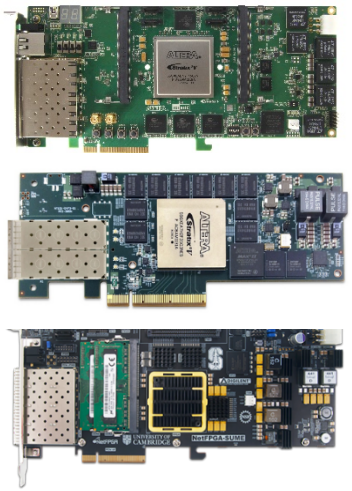

Experience – Towards a P4 FPGA-based SDN network Consensus as a Service (CAANS)

- Consensus protocols are the foundation for fault-tolerant systems
  - Ensures that a computation/group agrees on a value
  - E.g., OpenReplica, Ceph, Chubby
- Many distributed problems can be reduced to consensus
  - E.g., Atomic broadcast, atomic commit
- Any improvement in performance would have big impact

- Key Idea: Move Consensus into the Network

Conclusion

- *Chupja*: PHY covert channel
  - High-bandwidth, robust, and undetectable

- SoNIC Projects [NSDI '13]
  - P4FPGA / P4Paxos [arXiv'16; http://p4fpga.org]
  - Datacenter Time Protocol [SIGCOMM'16]
  - *Chupja:* Covert Channels [NSDI '14]
  - Understanding Burstiness [CISS '14]
  - MinProbe: Available bandwidth estimation [IMC '14]

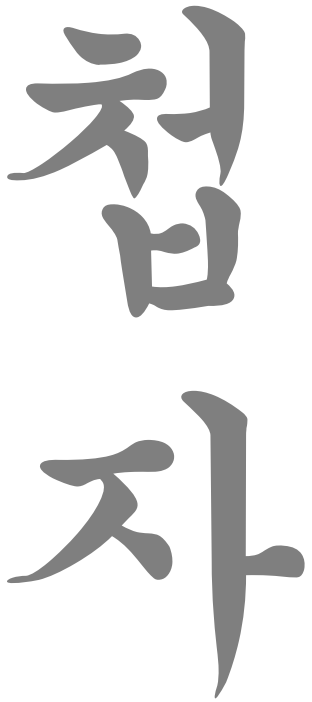

Contact: hweather@cs.cornell.edu
http://www.cs.cornell.edu/~hweather
Project website: http://sonic.cs.cornell.edu
Group website http://fireless.cs.cornell.edu

Thank you