Cornell University

Forensics in the SoNIC Project on Precise Realtime Software Access and Control of Wired Networks

Ki Suh Lee, Han Wang, Hakim Weatherspoon
Cornell University

International Workshop on Trustworthiness, Accountability, and Forensics in the Cloud (TAFC)
June 6, 2013

The Rise of Cloud Computing

- The promise of the Cloud
  - A computer utility; a commodity
  - Catalyst for technology economy
  - Revolutionizing for health care, financial systems, scientific research, and society

The Rise of Cloud Computing

- The promise of the Cloud
  - *ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.* NIST Cloud Definition

The Rise of Cloud Computing

- The promise of the Cloud
  - *ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.* NIST Cloud Definition

The Rise of Cloud Computing

- The promise of the Cloud
  - *ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.* NIST Cloud Definition

- How can we exploit the network for forensics, evidence, and accountability?
  - Public clouds: Bandwidth, availability
  - Private and hybrid clouds: exfiltration of data (covert channels)

Goal

Understand how to use the network to forensically account for and measure service level agreements in cloud

How to detect and/or prevent exfiltration of data from (private) clouds

Forensic Evidence via network interpacket delay

- Application
- Transport
- Network
- Data Link
- Physical

- Interpacket delay

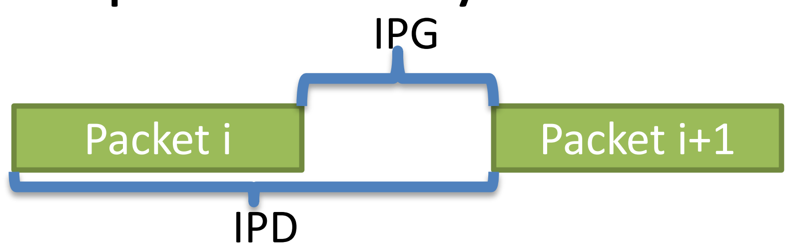

- Important metric for network forensic evidenc
  - Can be improved with access to the PHY

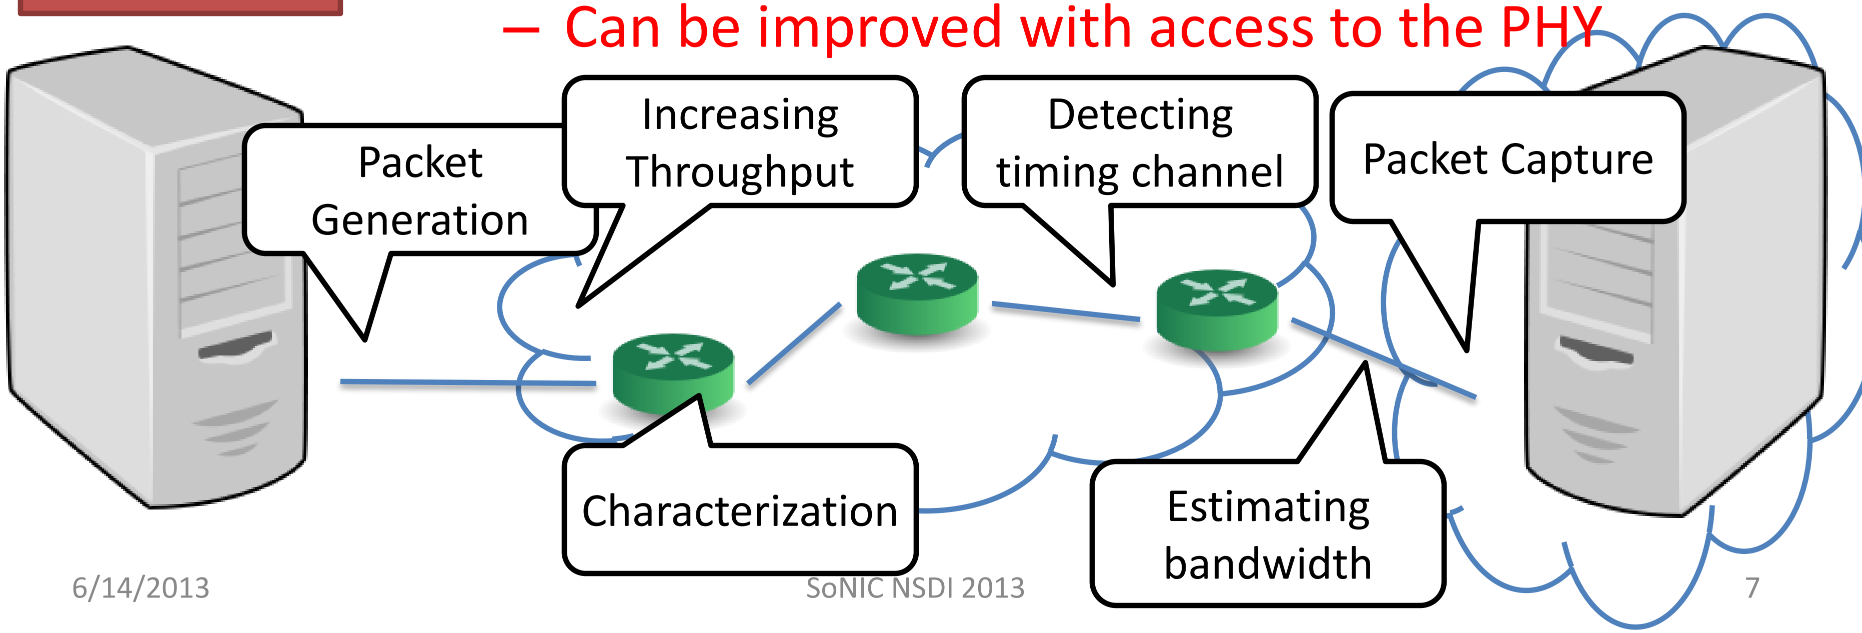

Forensic Evidence via network interpacket delay

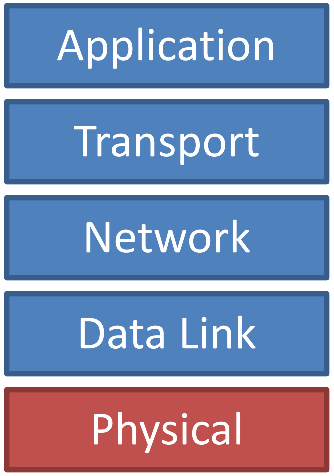

- Valuable information: Idle characters

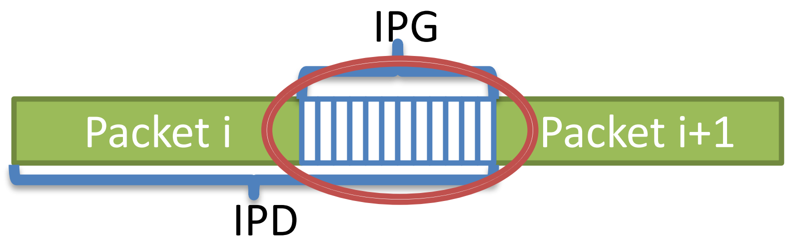

  - Can provide precise timing base for control
    - Each bit is ~97 *ps* wide

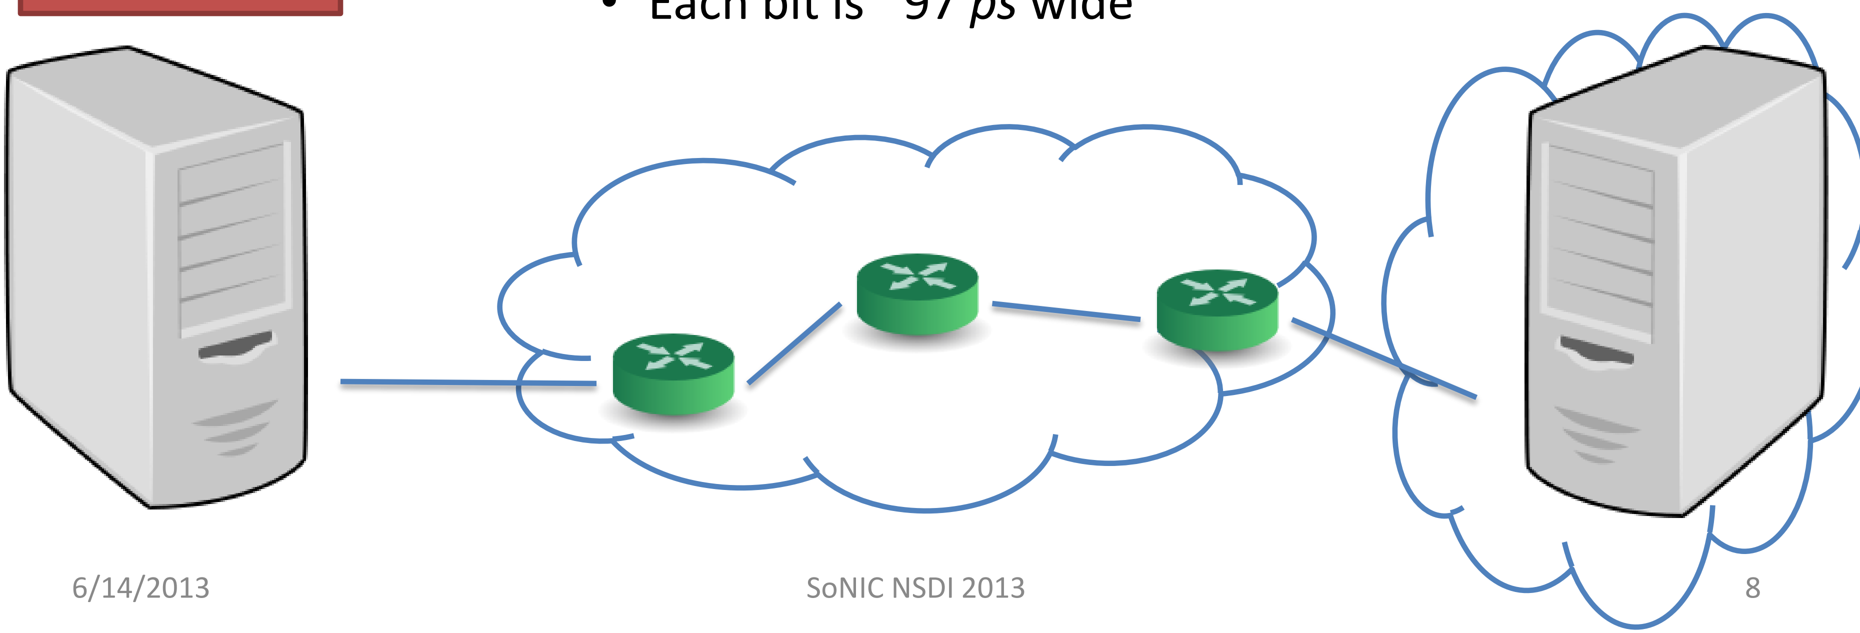

Forensic Evidence via network interpacket delay

- Application
- Transport
- Network
- Data Link
- Physical

- Valuable informa... rs
  - Ca... ming base for control
    - Each bit is ~97 *ps* wide

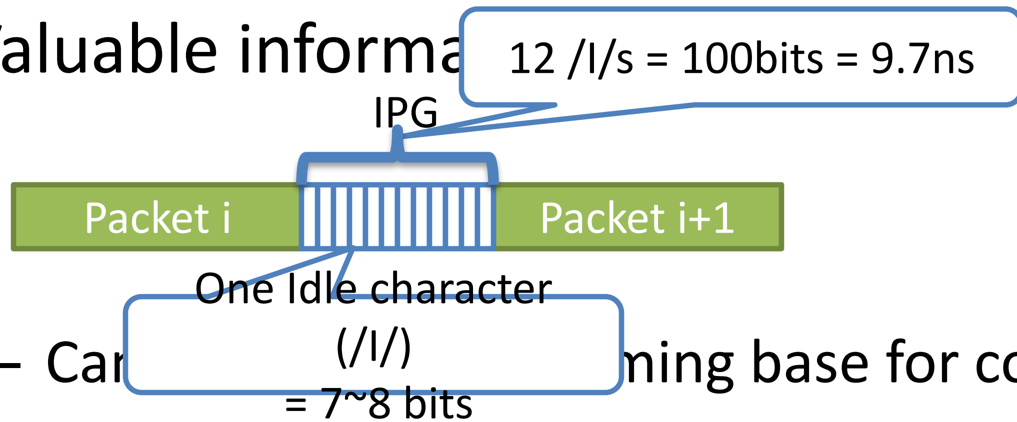

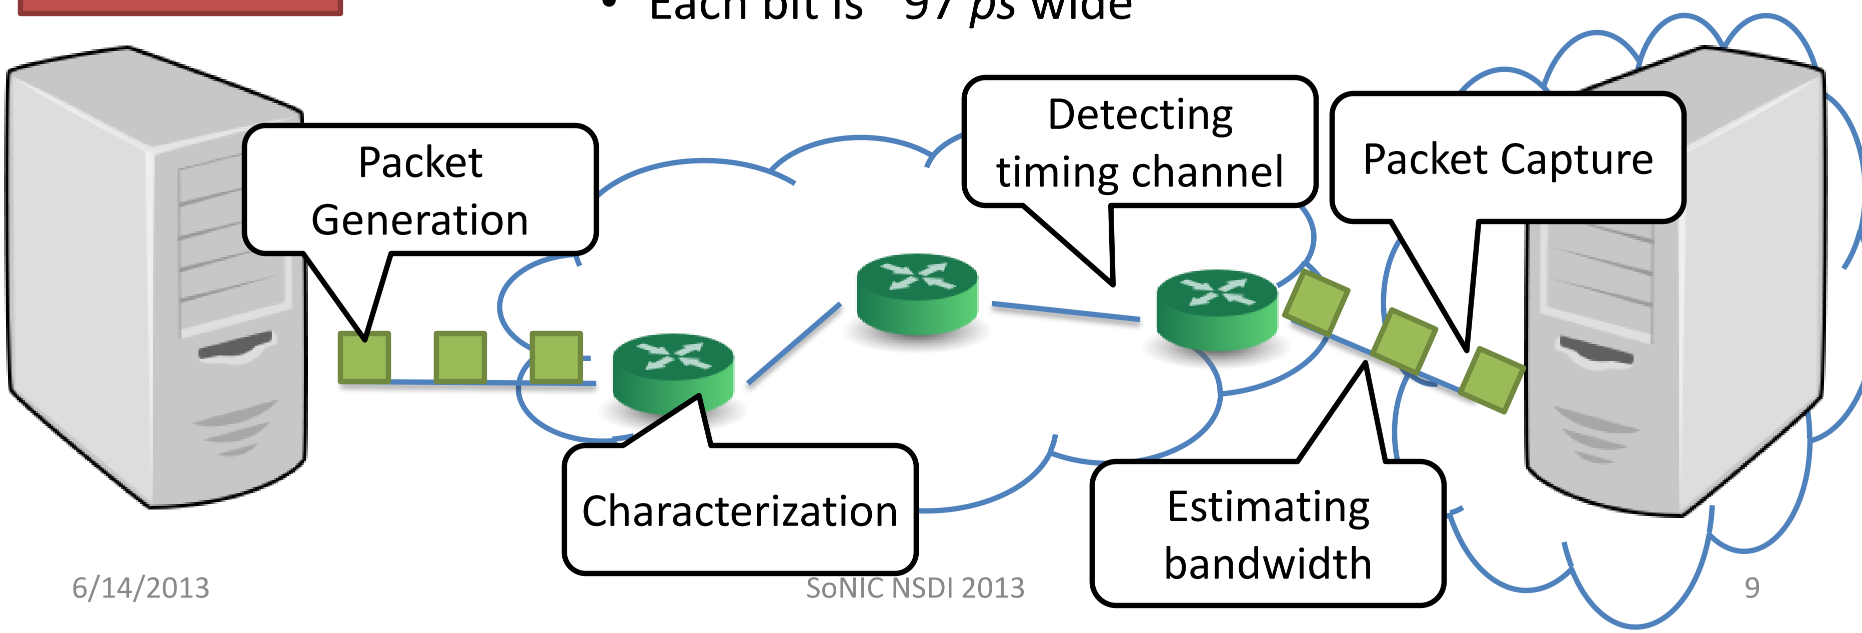

Forensic Evidence via network interpacket delay

- Application
- Transport
- Network
- Data Link
- Physical

- Valuable information in PHY: Idle characters

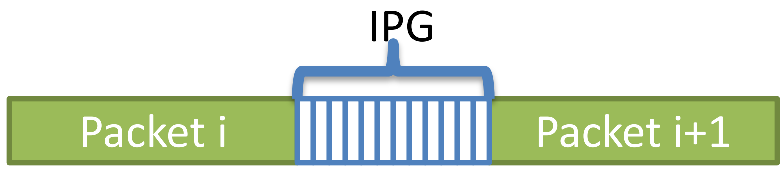

- Issue1: The PHY is simply a black box
  - No interface from NIC or OS
  - Valuable information is invisible (discarded)

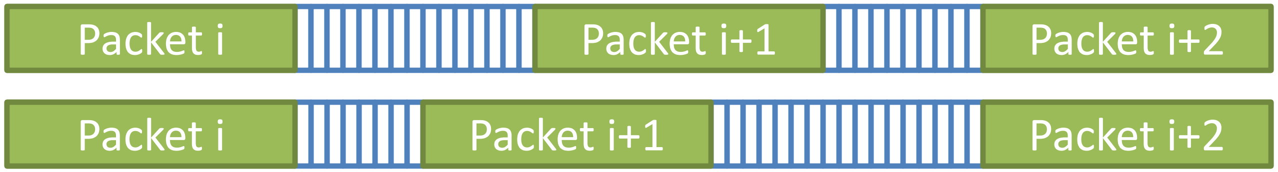

- Issue2: Limited access to hardware

Forensic Evidence via network interpacket delay

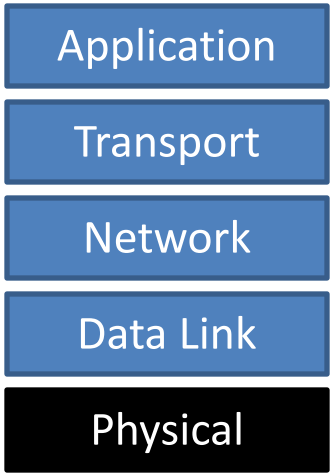

- Goal: Control *every* bit in *software* in *realtime*

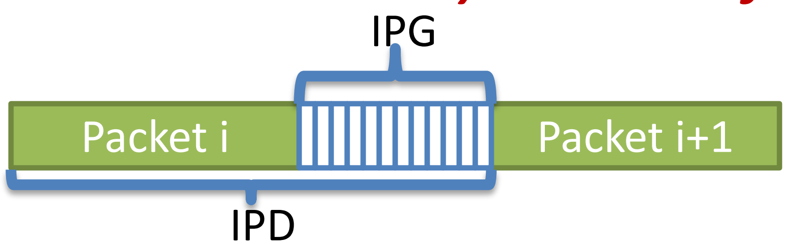

  – Enable research on PHY covert challenge

- Challenge
  – Requires unprecedented software access to the PHY

SoNIC: Software-defined Network Interface Card

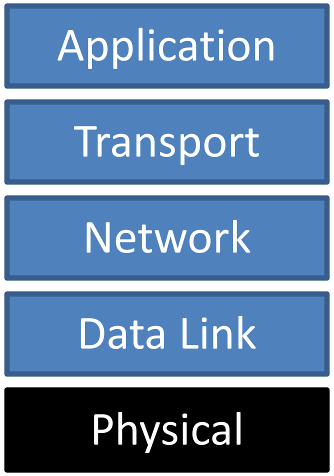

- Implements the PHY in software

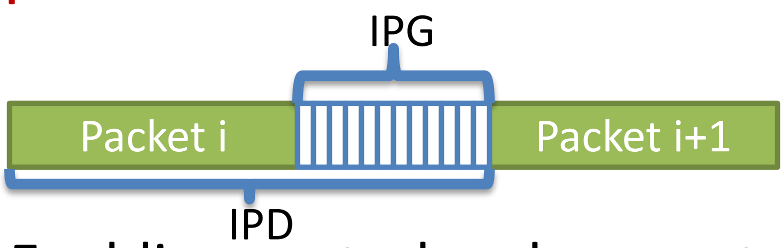

  - Enabling control and access to every bit in realtime
  - With commodity components
  - Thus, enabling novel network research

SoNIC: Precise Realtime Software Access and Control of Wired Networks, Ki Suh Lee, Han Wang and Hakim Weatherspoon, Appears in NSDI, April 2013

Outline

- Introduction
- Examples of Forensic Evidence
  - Available bandwidth estimation
  - PHY Covert Timing Channel
- SoNIC: Software-defined Network Interface Card
- Concluding Remarks

SoNIC

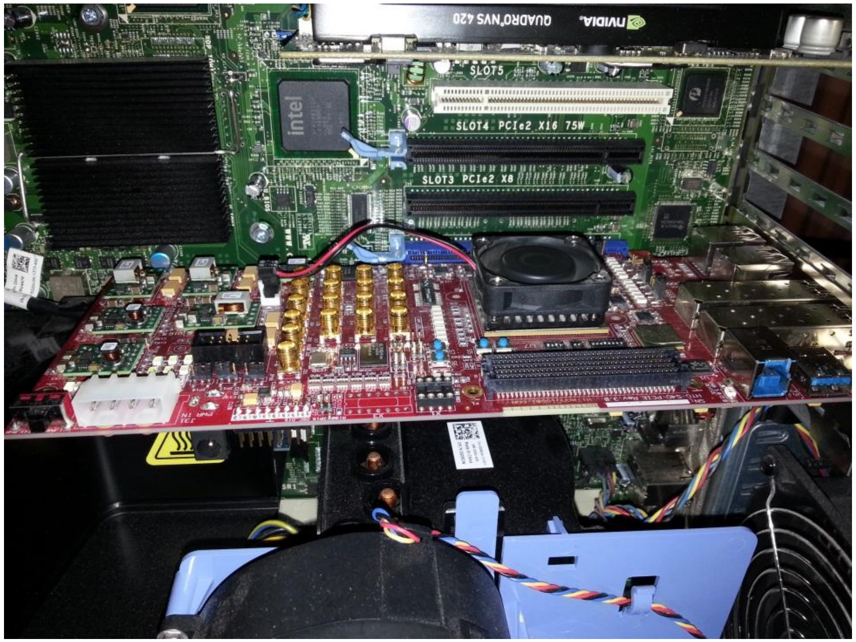

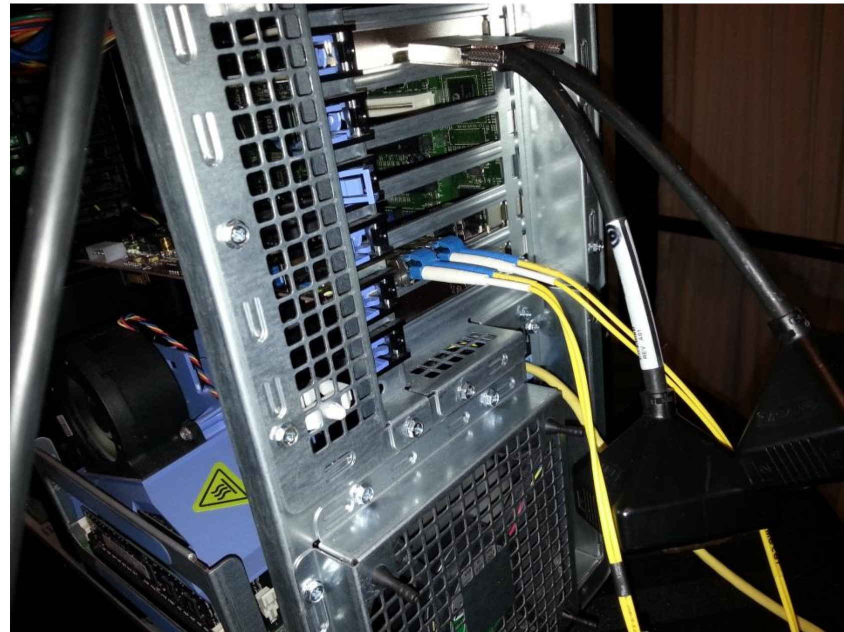

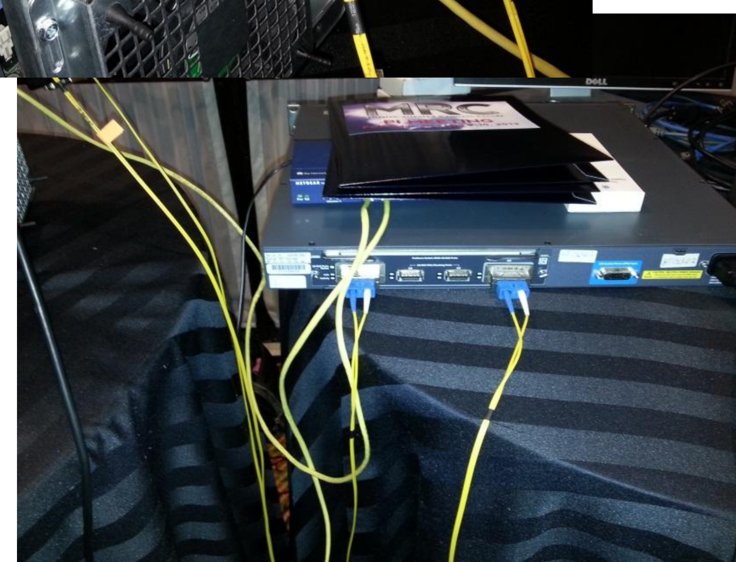

Forensic Evidence: Bandwidth Estimation

- Estimate available bandwidth
  - Traffic sent, packet trains:

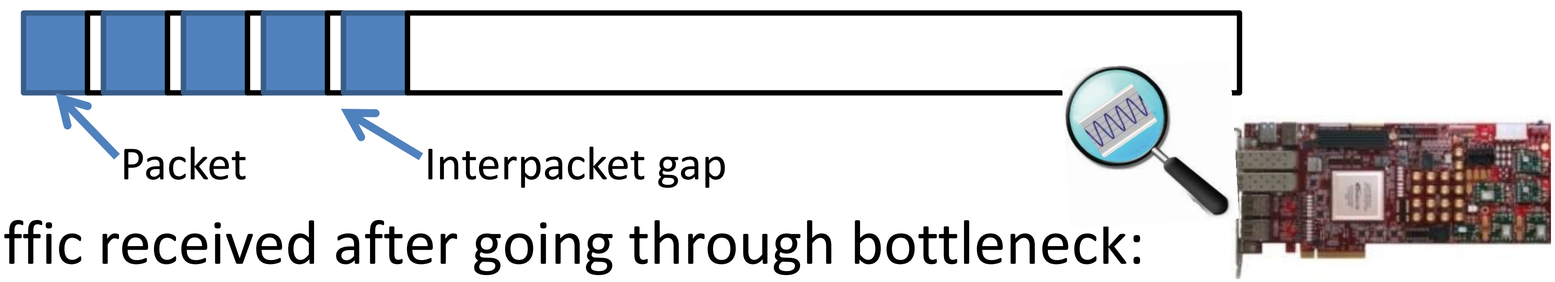

  - Traffic received after going through bottleneck:

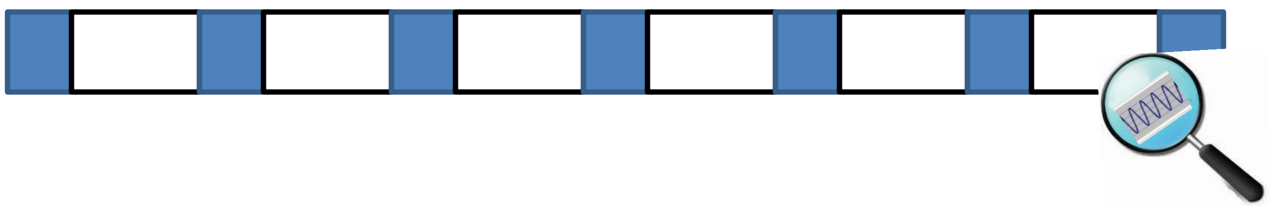

- Accurate available bandwidth estimation requires PHY
- Inter-packet gaps are invisible to higher layers, but not SoNIC

Outline

- Introduction
- Examples of Forensic Evidence
  - Available bandwidth estimation
  - **PHY Covert Timing Channel**
- SoNIC: Software-defined Network Interface Card
- Concluding Remarks

Forensic Evidence: Covert Timing Channel

- Embedding signals into interpacket gaps.
  - Large gap: ‘1’
  - Small gap: ‘0’

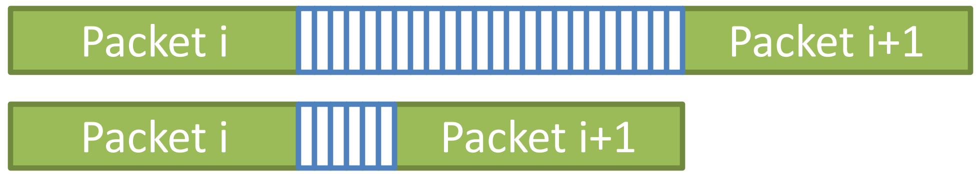

- Covert timing channel by modulating IPGs at 100ns
    - Overt channel at 3 Gbps
    - Covert channel at 250 kbps
    - Over 4-hops with < 1% BER

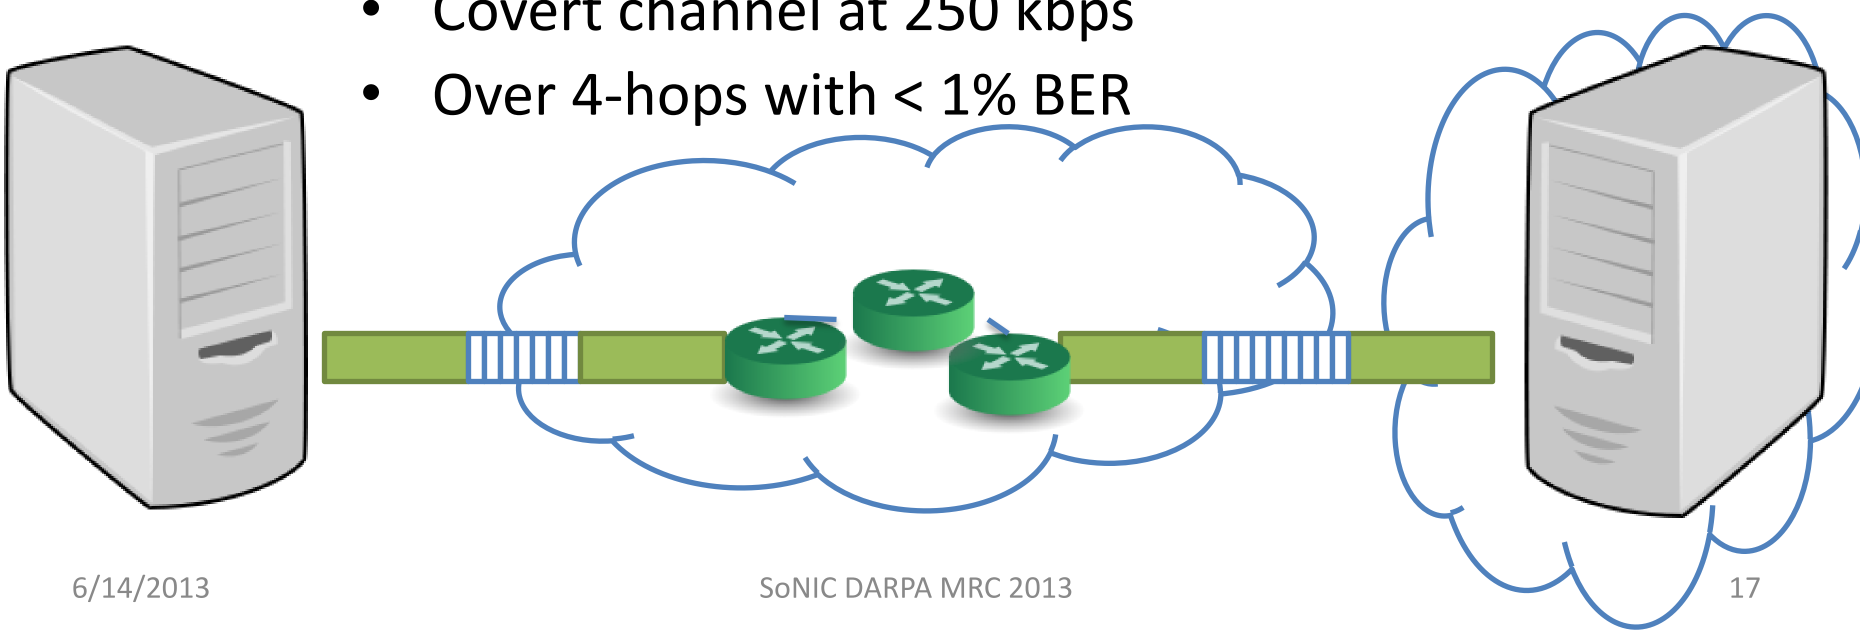

Forensic Evidence: Covert Timing Channel

- *Modulating IPGs at 100ns scale (=128 /I/s), over 4 hops*

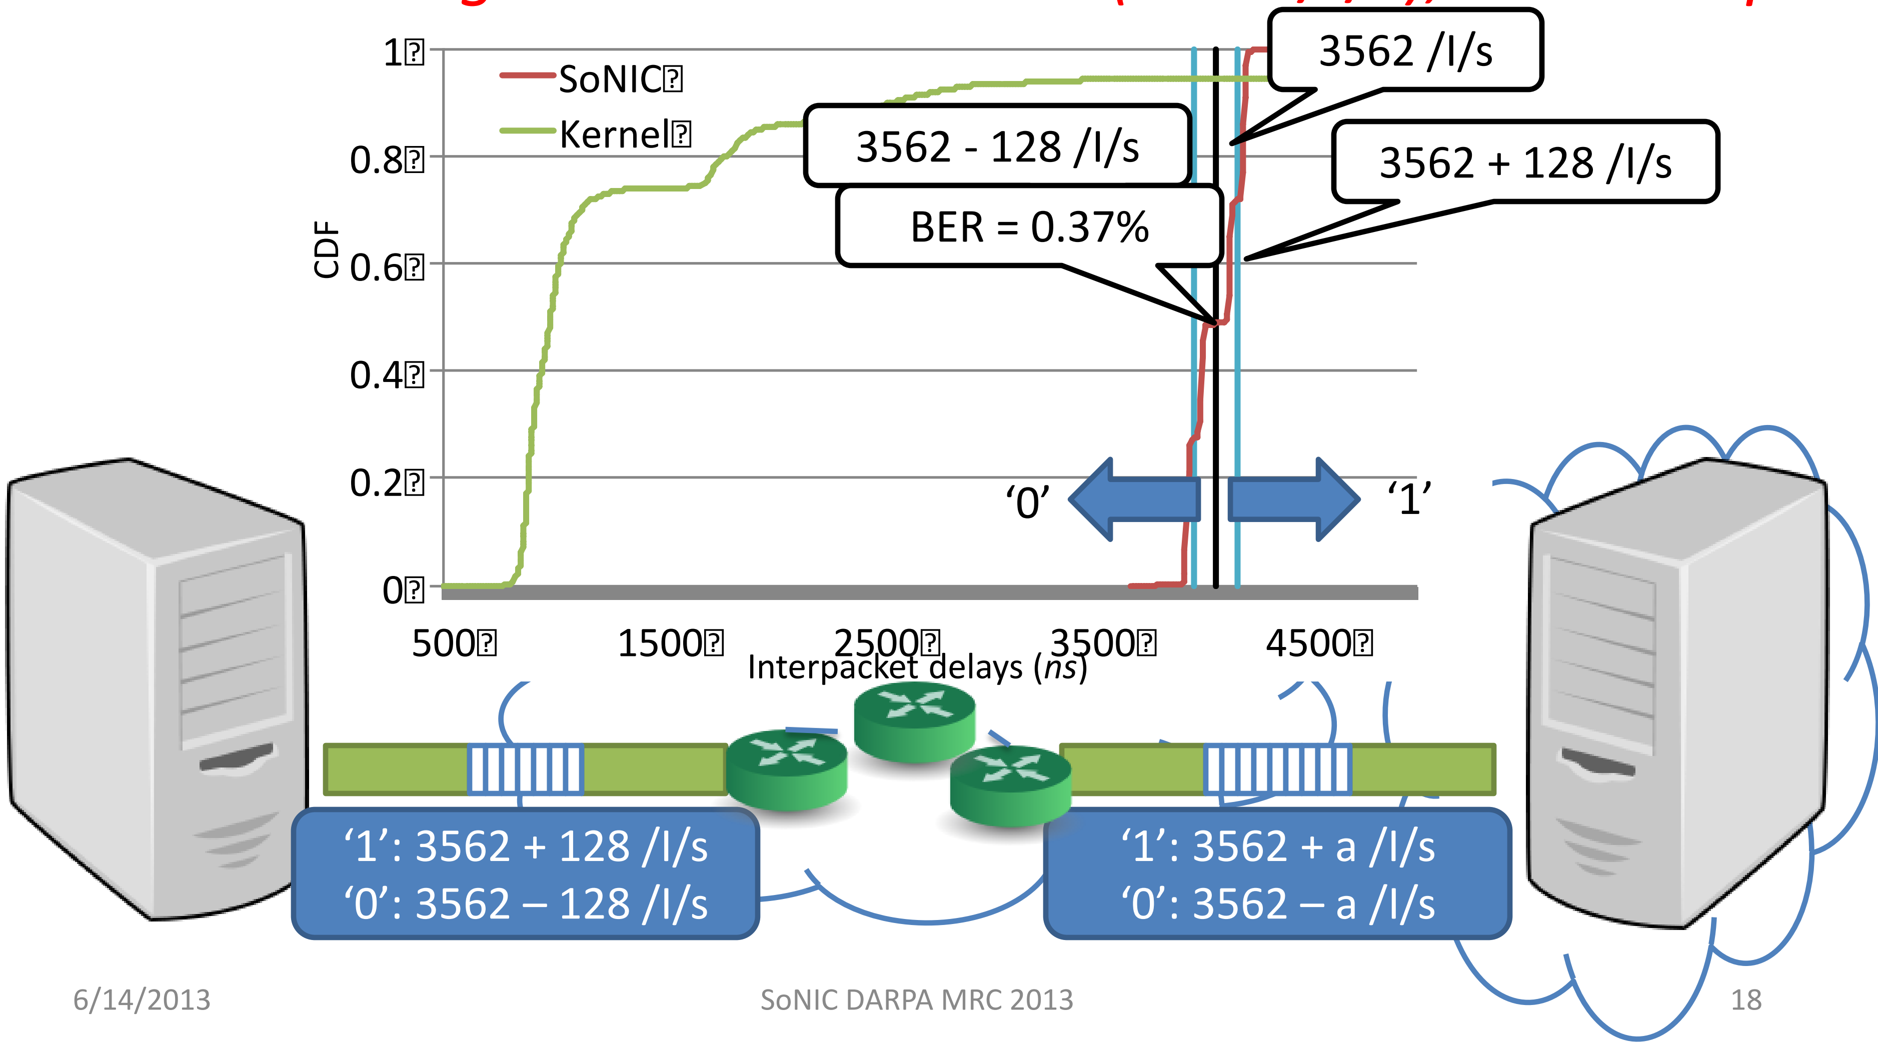

Forensic Evidence: Covert Timing Channel

- *Prevent Covert Timing Channels?*

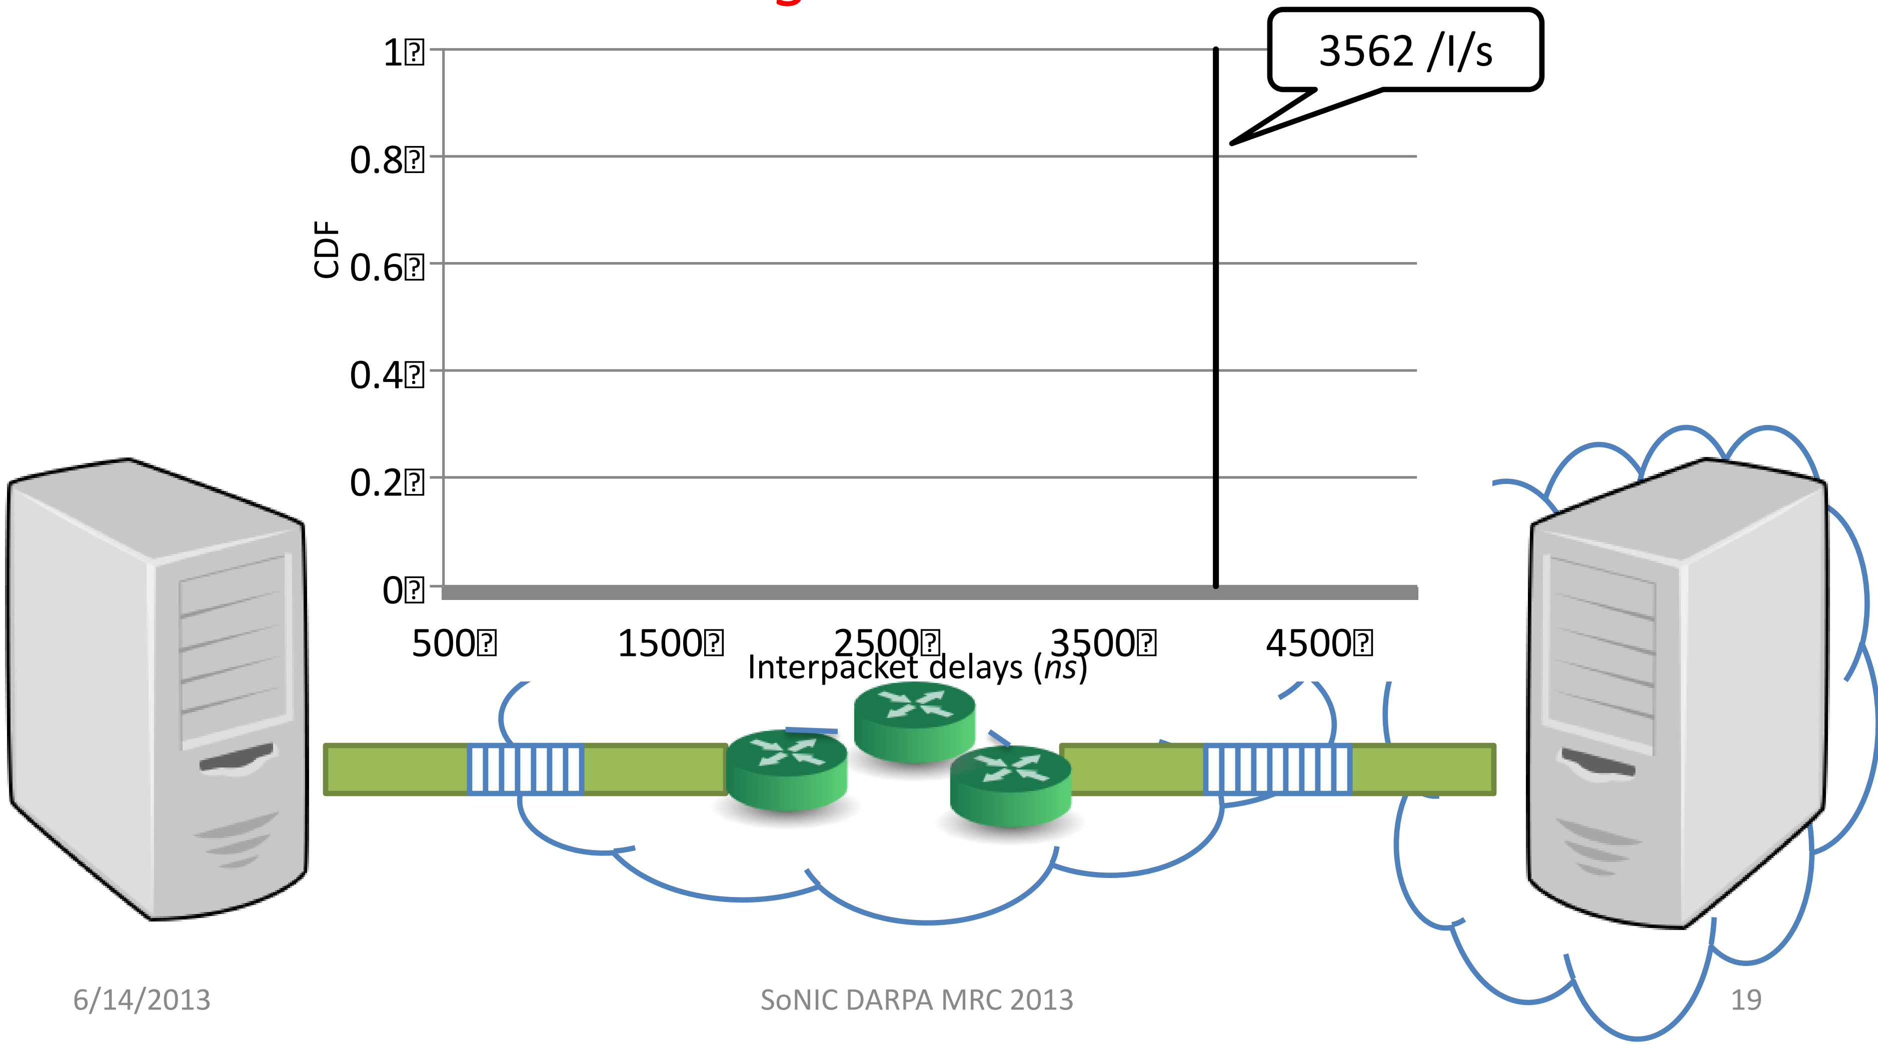

Forensic Evidence: Covert Timing Channel

- Router/ Switch Signatures
  - Different Routers and switches have different response function.
  - Improve simulation model of switches and routers.
  - Detect switch and router model in real network.

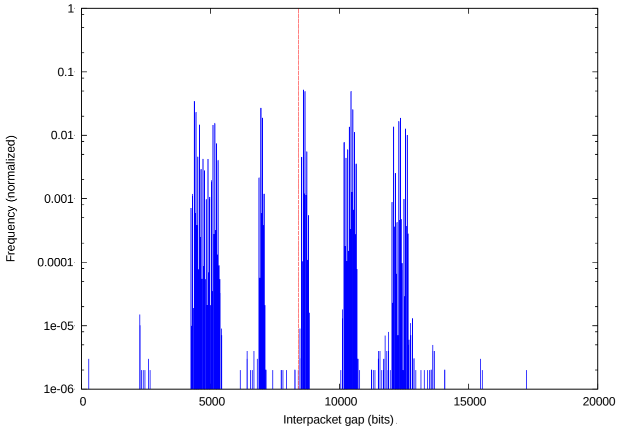

Cisco 4948

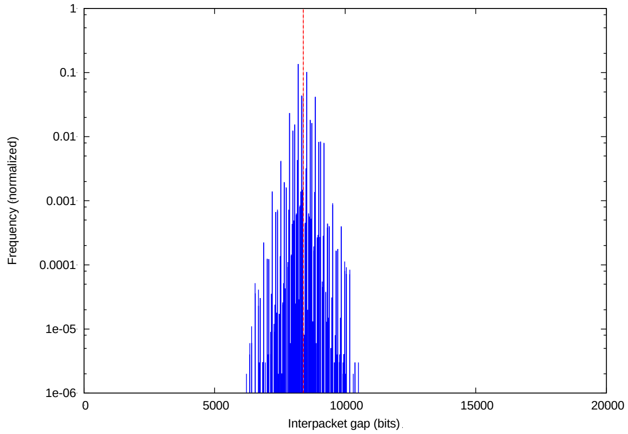

Cisco 6509

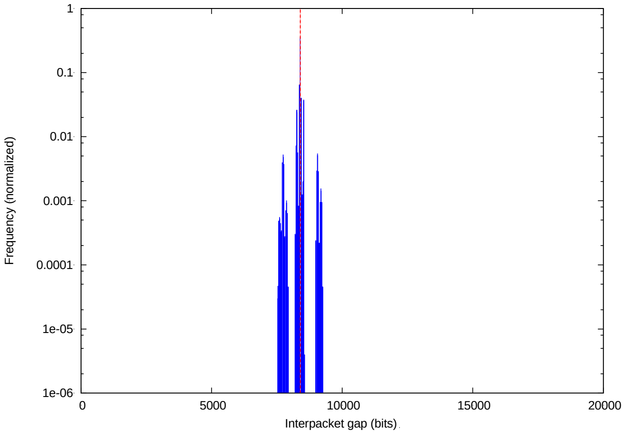

IBM BNT G8264R

1500 byte packets @ 6Gbps

Outline

- Introduction
- Demo: PHY Covert Timing Channel
- SoNIC: Software-defined Network Interface Card
- Concluding Remarks

10GbE Network Stack

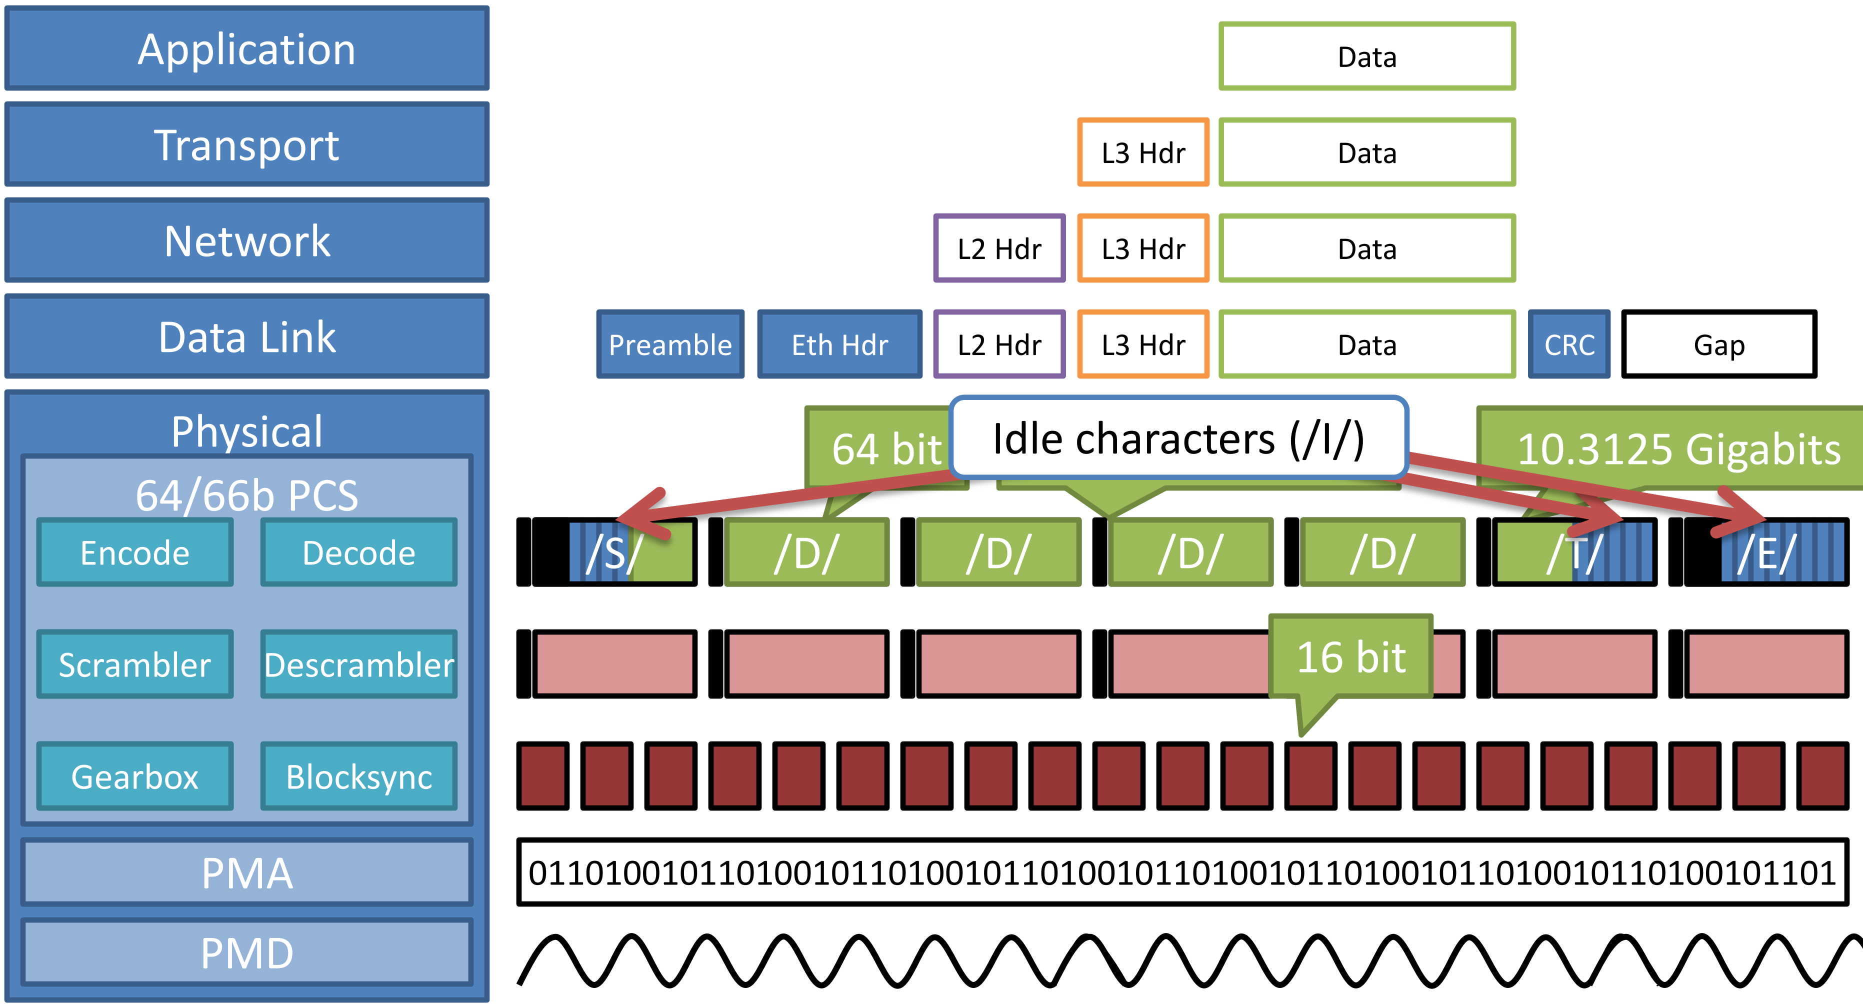

10GbE Network Stack

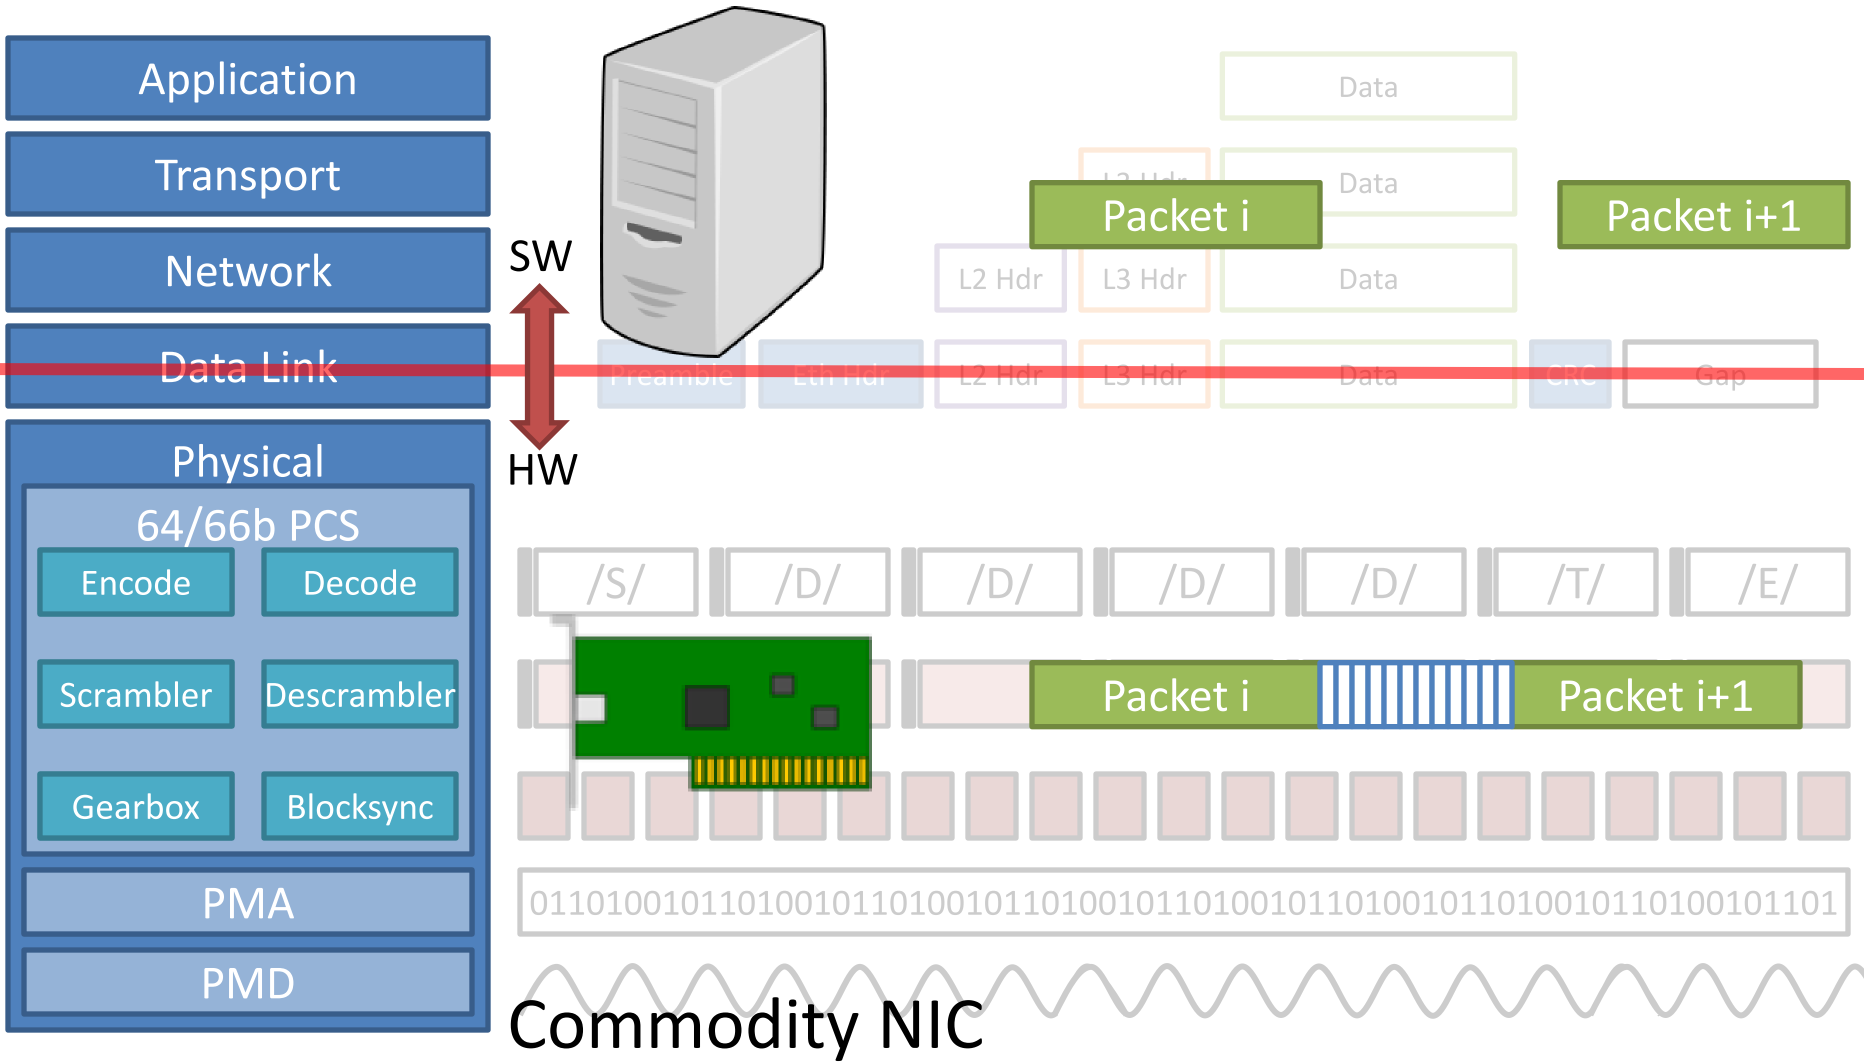

10GbE Network Stack

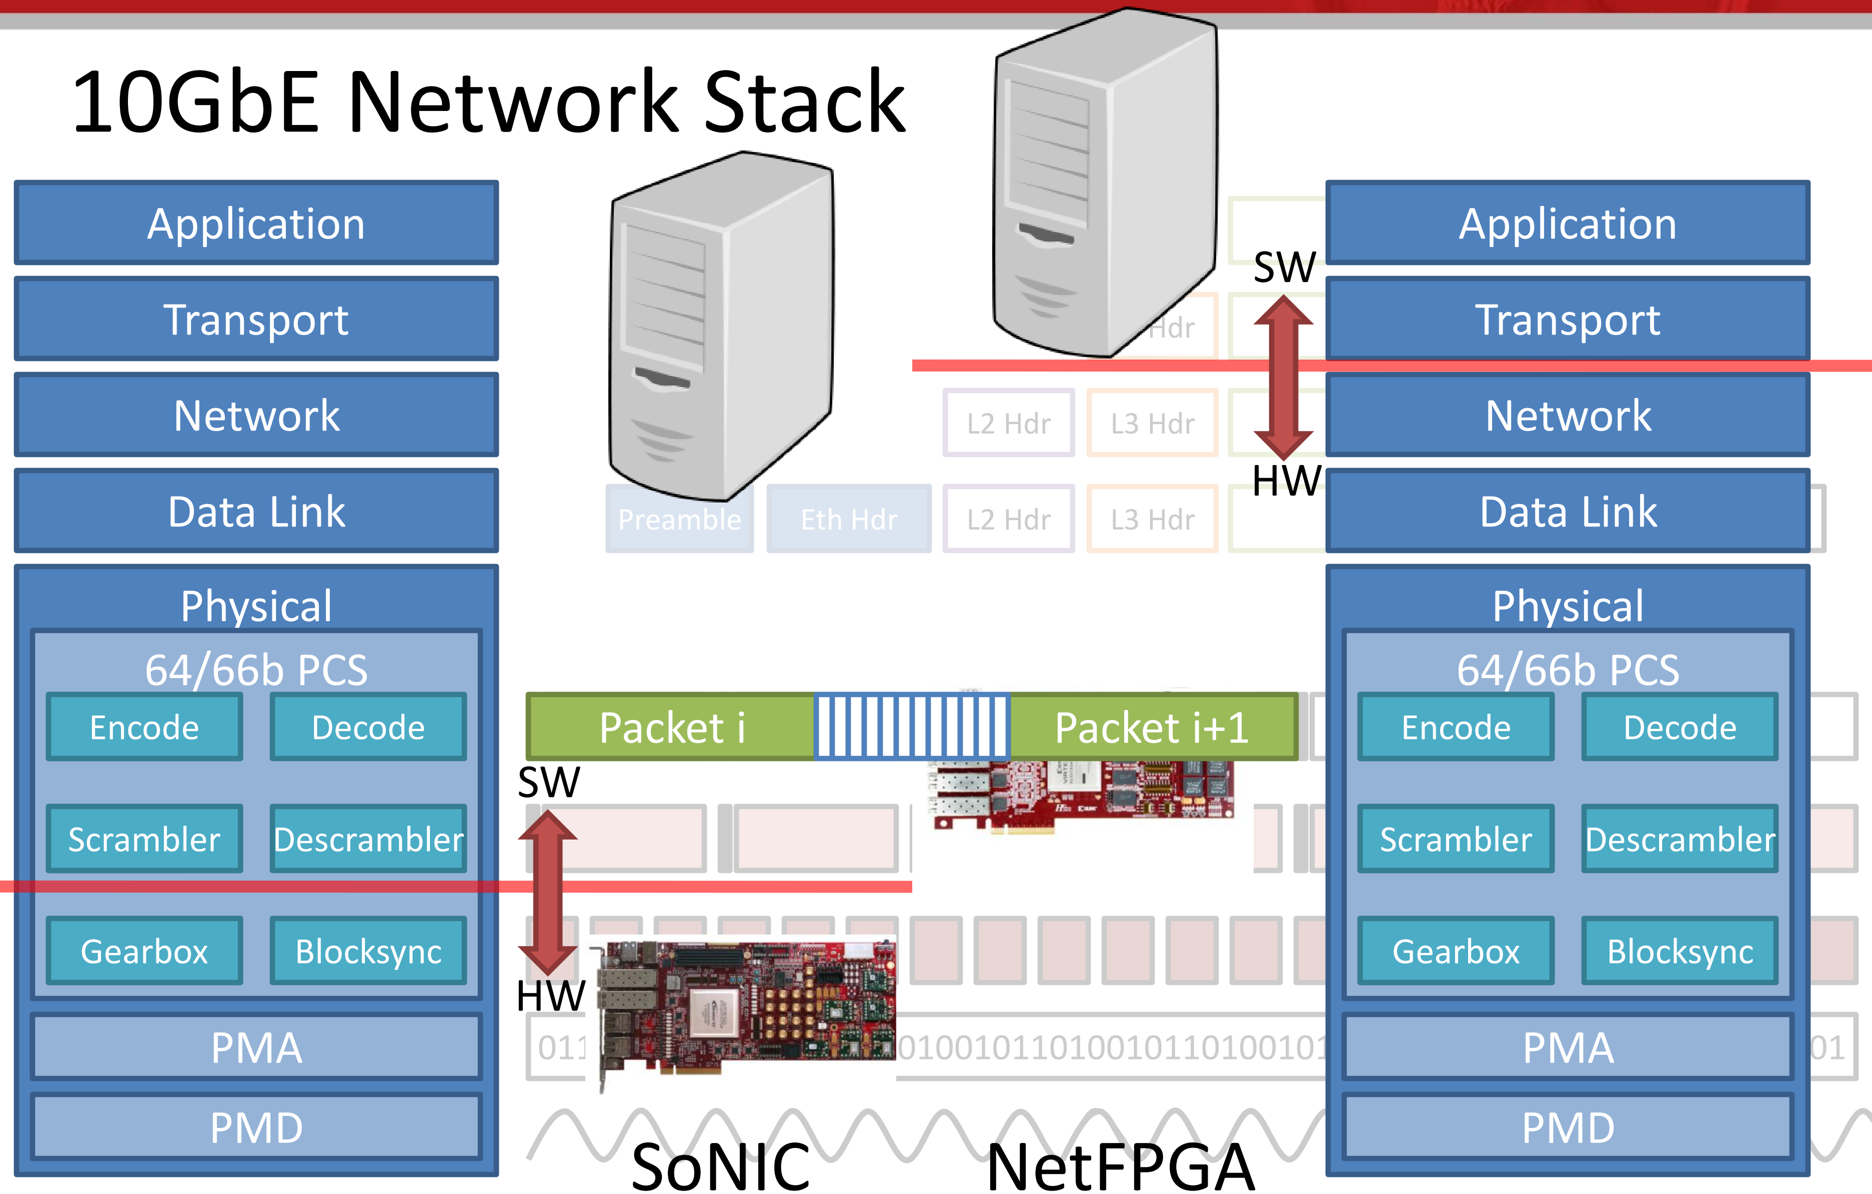

SoNIC Design

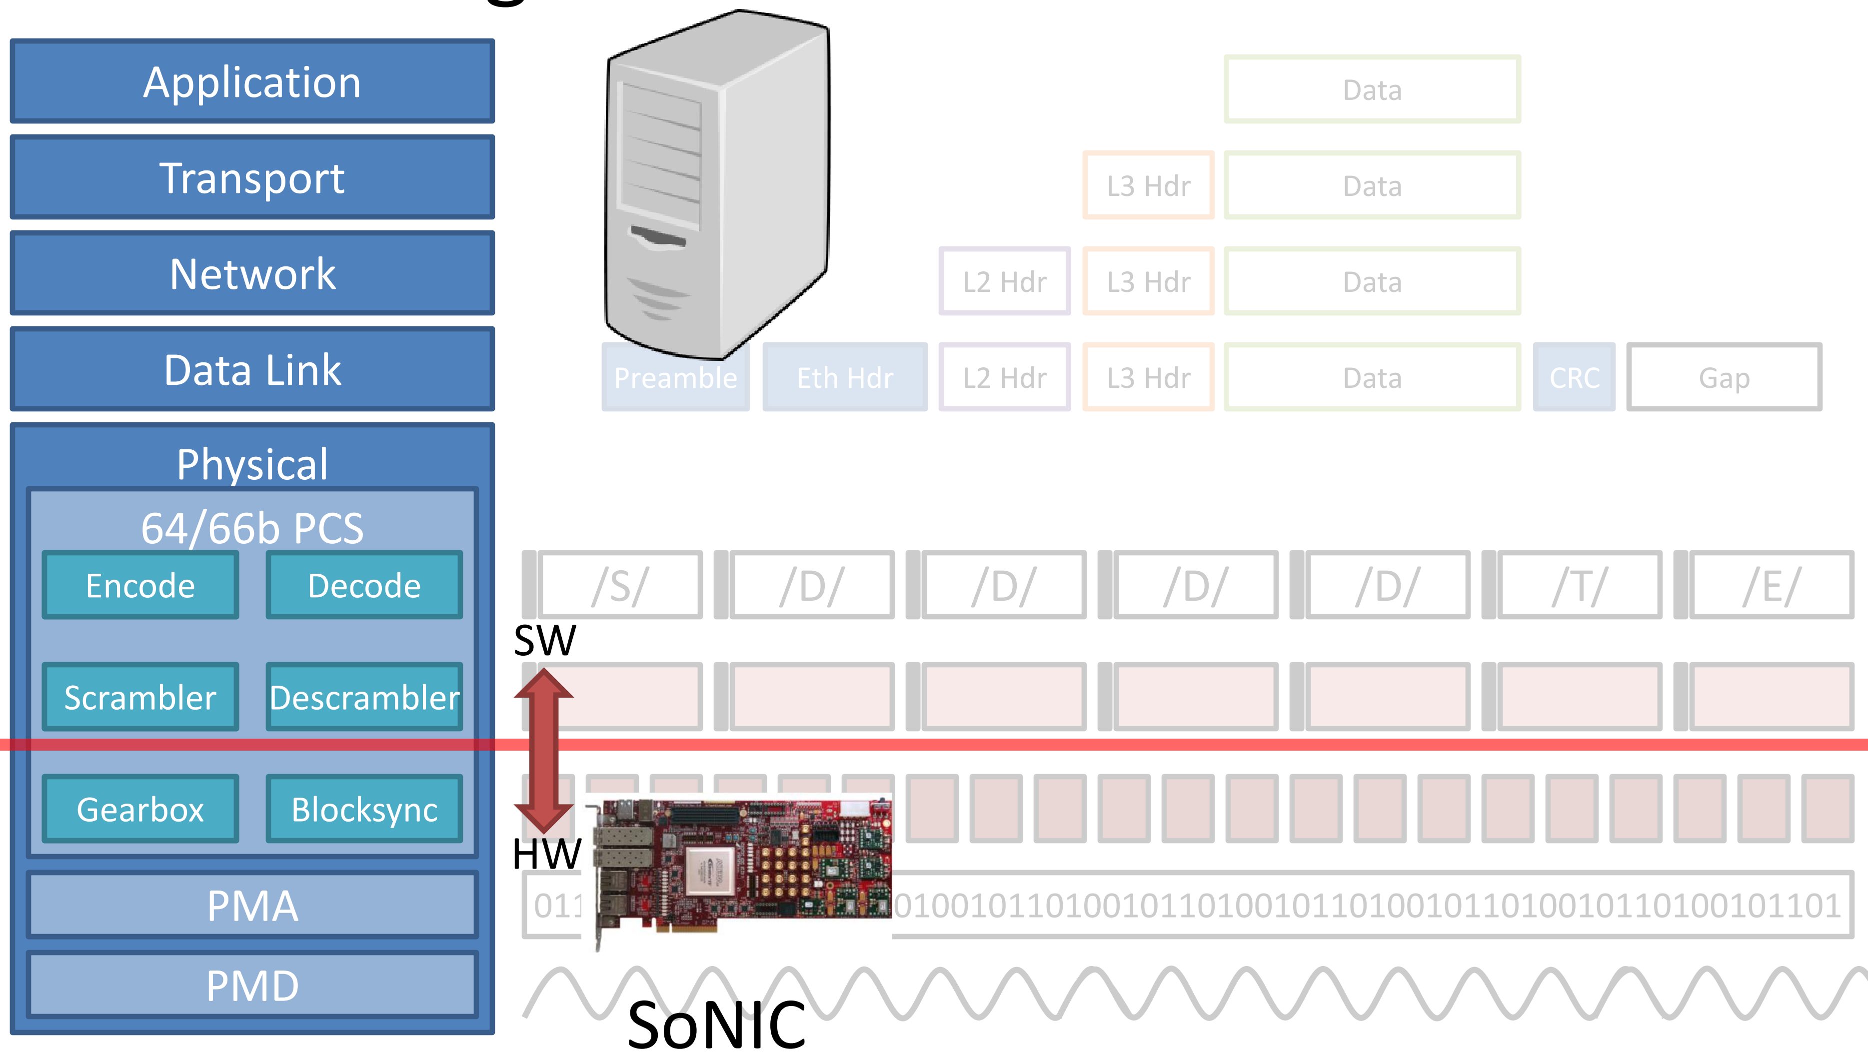

SoNIC Design and Architecture

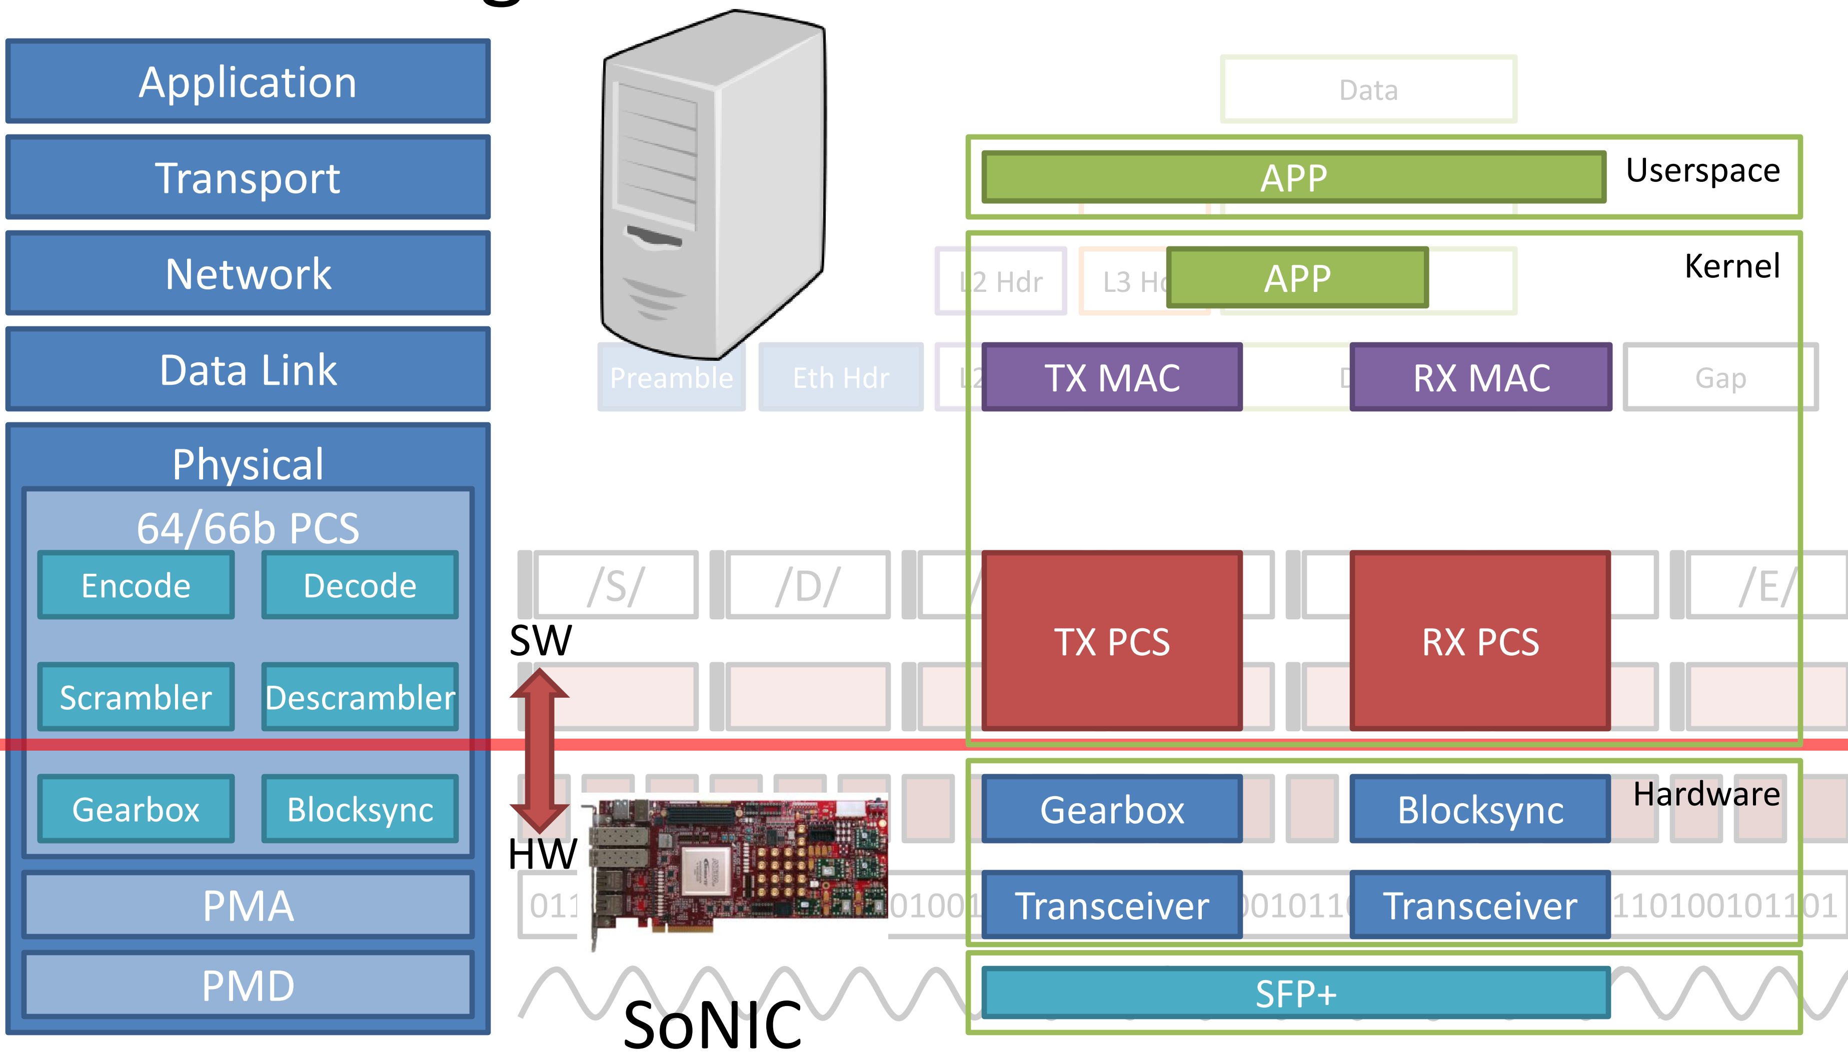

SoNIC Design: Interface and Control

- Hardware control: *ioctl* syscall
- I/O : character device interface
- Sample C code for packet generation and capture

```
#include "sonic.h"

struct sonic_pkt_gen_info info = {
.mode = 0,
.pkt_num = 1000000000UL,
.pkt_len = 1518,
.mac_src = "00:11:22:33:44:55",
.mac_dst = "aa:bb:cc:dd:ee:ff",
.ip_src = "192.168.0.1",
.ip_dst = "192.168.0.2",
.port_src = 5000,
.port_dst = 5000,
.idle = 12,
};

/* OPEN DEVICE*/
fd1 = open(SONIC_CONTROL_PATH, O_RDWR);
fd2 = open(SONIC_PORT1_PATH, O_RDONLY);
/* CONFIG SONIC CARD FOR PACKET GEN*/
ioctl(fd1, SONIC_IOC_RESET)
ioctl(fd1, SONIC_IOC_SET_MODE, PKT_GEN_CAP)
ioctl(fd1, SONIC_IOC_PORT0_INFO_SET, &info)

/* START EXPERIMENT*/
ioctl(fd1, SONIC_IOC_START)
// wait till experiment finishes
ioctl(fd1, SONIC_IOC_STOP)

/* CAPTURE PACKET */
while ((ret = read(fd2, buf, 65536)) > 0) {
// process data
}

close(fd1);
close(fd2);
```

Contributions

- Network Research
  - Unprecedented access to the PHY with commodity hardware
  - A platform for cross-network-layer research
  - Can improve network research applications
- Engineering
  - Precise control of interpacket gaps (delays)
  - Design and implementation of the PHY in software
  - Novel scalable hardware design
  - Optimizations / Parallelism
- Status
  - Measurements in large scale: DCN, GENI, 40 GbE

Concluding Remarks

- The network is at the center of the cloud
  - SoNIC gives precise realtime software access and control of the network
  - Necessary for forensics, evidence, and accountability of network/cloud
- Network is useful to validate SLAs
  - Accurate bandwidth estimation
  - Characterize/profile/fingerprint network components
- Need to understand entire network stack to protect data
  - Demonstrate: Covert Timing Channel
  - 4 hops, 250kbps, less than 1% BER
- Status
  - SoNIC in large scale: DURIP, GENI, 40 GbE
  - http://sonic.cs.cornell.edu
  - SoNIC is available Open Source.

My Contributions

- Cloud Networking
  - SoNIC in NSDI 2013
  - Wireless DC in ANCS 2012 (best paper) and NetSlice in ANCS 2012
  - Bifocals in IMC 2010 and DSN 2010
  - Maelstrom in ToN 2011 and NSDI 2008
  - Chaired Tudor Marian's PhD 2010 (now at Google)
- Cloud Computation & Vendor Lock-in
  - Plug into the Supercloud in IEEE Internet Computing-2013
  - Supercloud/Xen-Blanket in EuroSys-2012 and HotCloud-2011
  - Overdriver in VEE-2011
  - Chaired Dan William's PhD 2012 (now at IBM)
- Cloud Storage
  - Gecko in FAST 2013 / HotStorage 2012
  - RACS in SOCC-2010
  - SMFS in FAST 2009
  - Antiquity in EuroSys 2007 / NSDI 2006
  - Chaired Lakshmi Ganesh's PhD 2011 (now at UT Austin)

Thank you!

http://sonic.cs.cornell.edu