Revisiting Traffic Anomaly Detection using Software Defined Networking

Syed Akbar Mehdi, Junaid Khalid, Syed Ali Khayam

School of Electrical Engineering and Computer Science,
National University of Sciences and Technology, Pakistan

Home Network Security

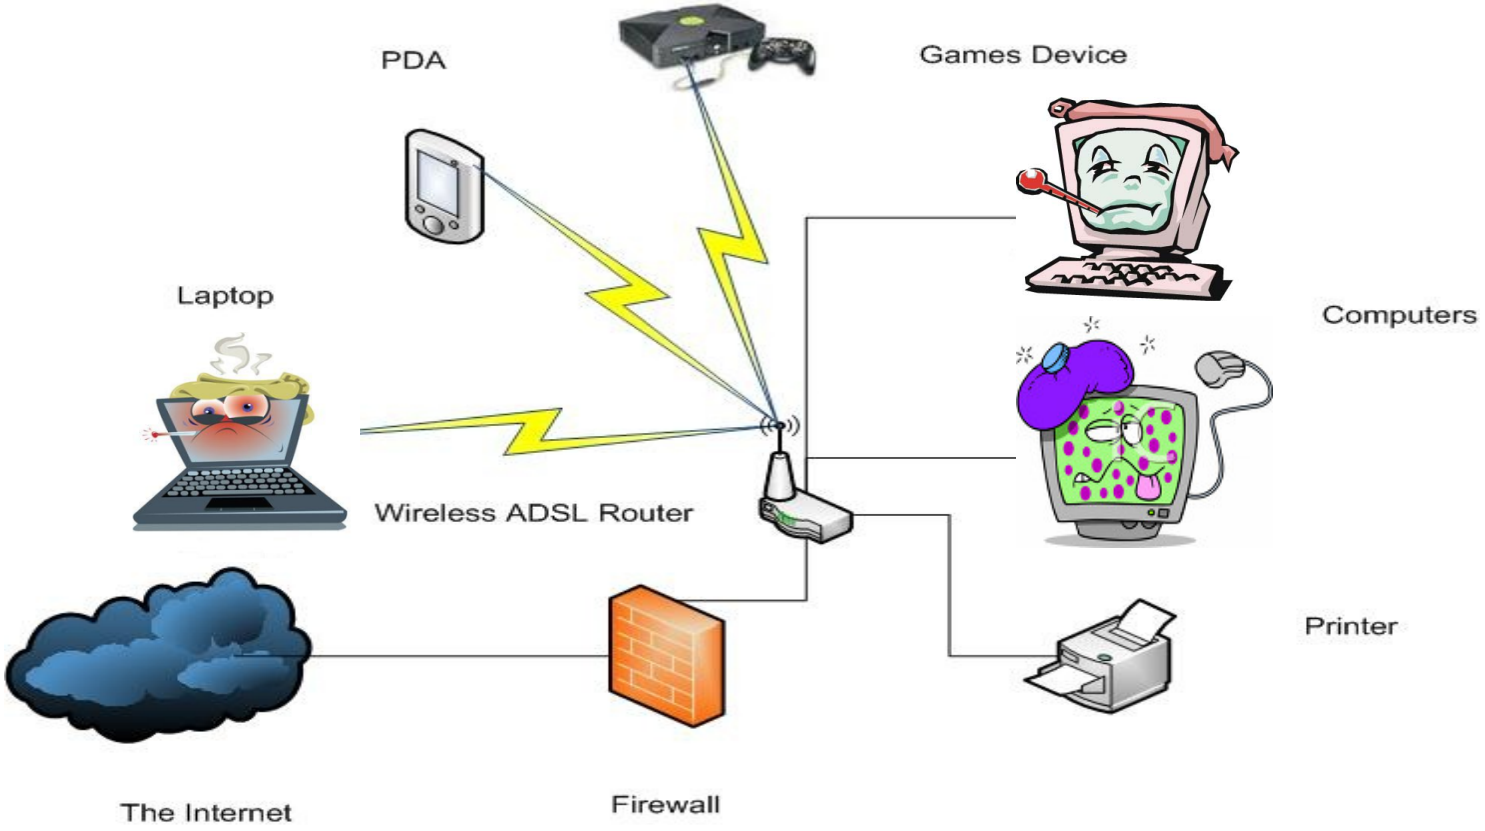

- Usually no expert administrator.
- Security problems (e.g. malware infections) can have consequences.
  - Personal Loss (e.g. stolen credit cards and passwords)
  - Internet-wide security issues (e.g. DDoS, Spam Email)

Home Network Security

- Infected hosts are usually part of botnets.
- E.g. Torpig Botnet*

Network Speed	Unique infected hosts
Cable/DSL	50,535
Dial-up	9,923
Corporate	17,217
Unknown	105,125
	182,800

DDoS potential: 17Gbps

- Conficker botnet -- More than 10 million infected worldwide.

*Source: "*Your botnet is my botnet: Analysis of a botnet takeover*" by Stone-Gross et.al. (ACM CCS 2009)

Solutions?

- Detect in the network core (e.g. ISP)
- Detect at the network edge (e.g. home network)

Solutions?

- Detect in the network core (e.g. ISP)

- Detect at the network edge (e.g. home network)

Detect at the ISP ?

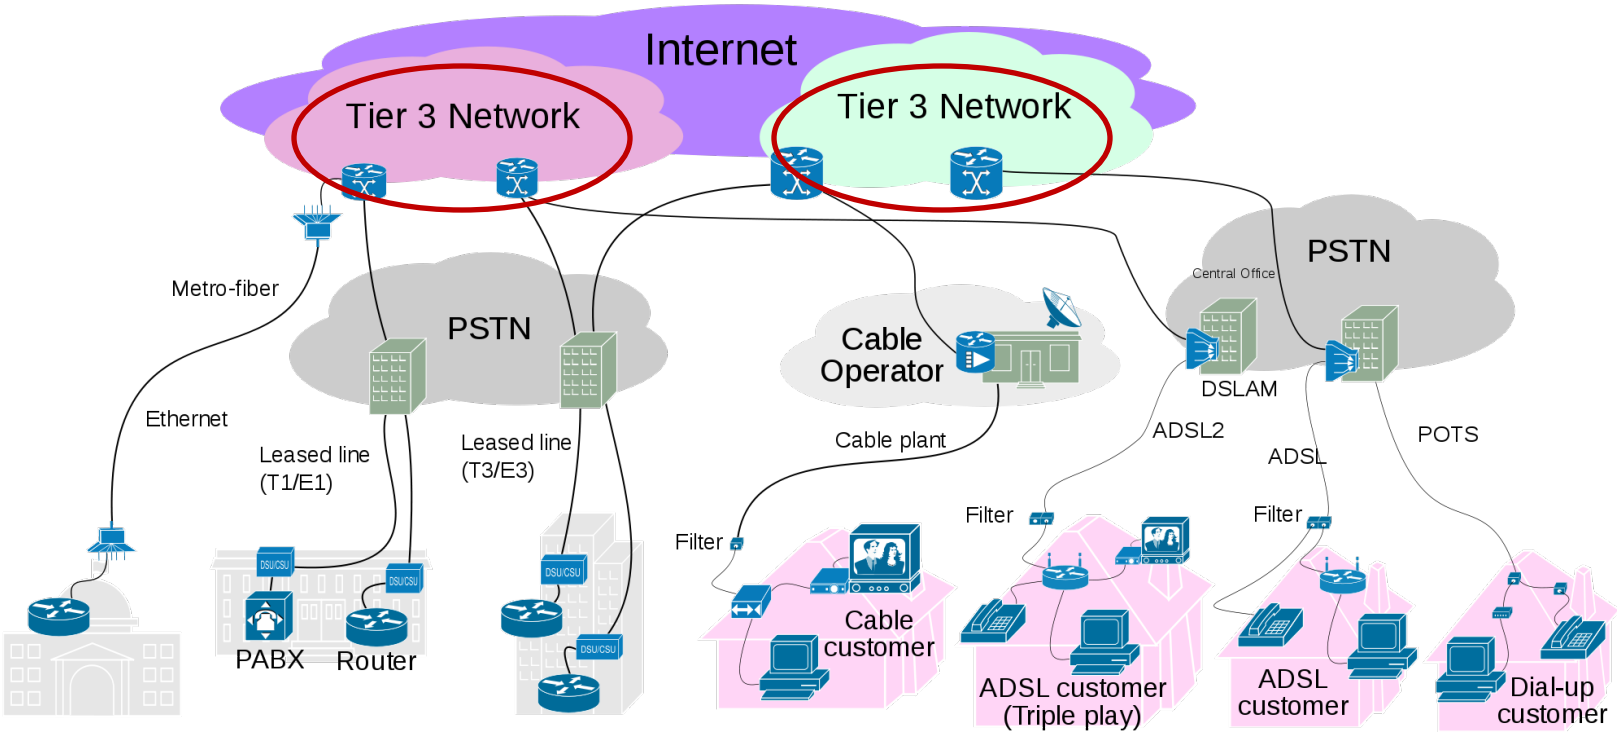

- Problem: **Scalability**
- Real-time detection requires *in-line processing*.
- ISPs are hesitant to do complicated processing in forwarding path of network traffic.
  - Rising traffic rates and volumes
  - Stalled processor clock speeds.

Image courtesy Wikipedia

Detect at the ISP ?

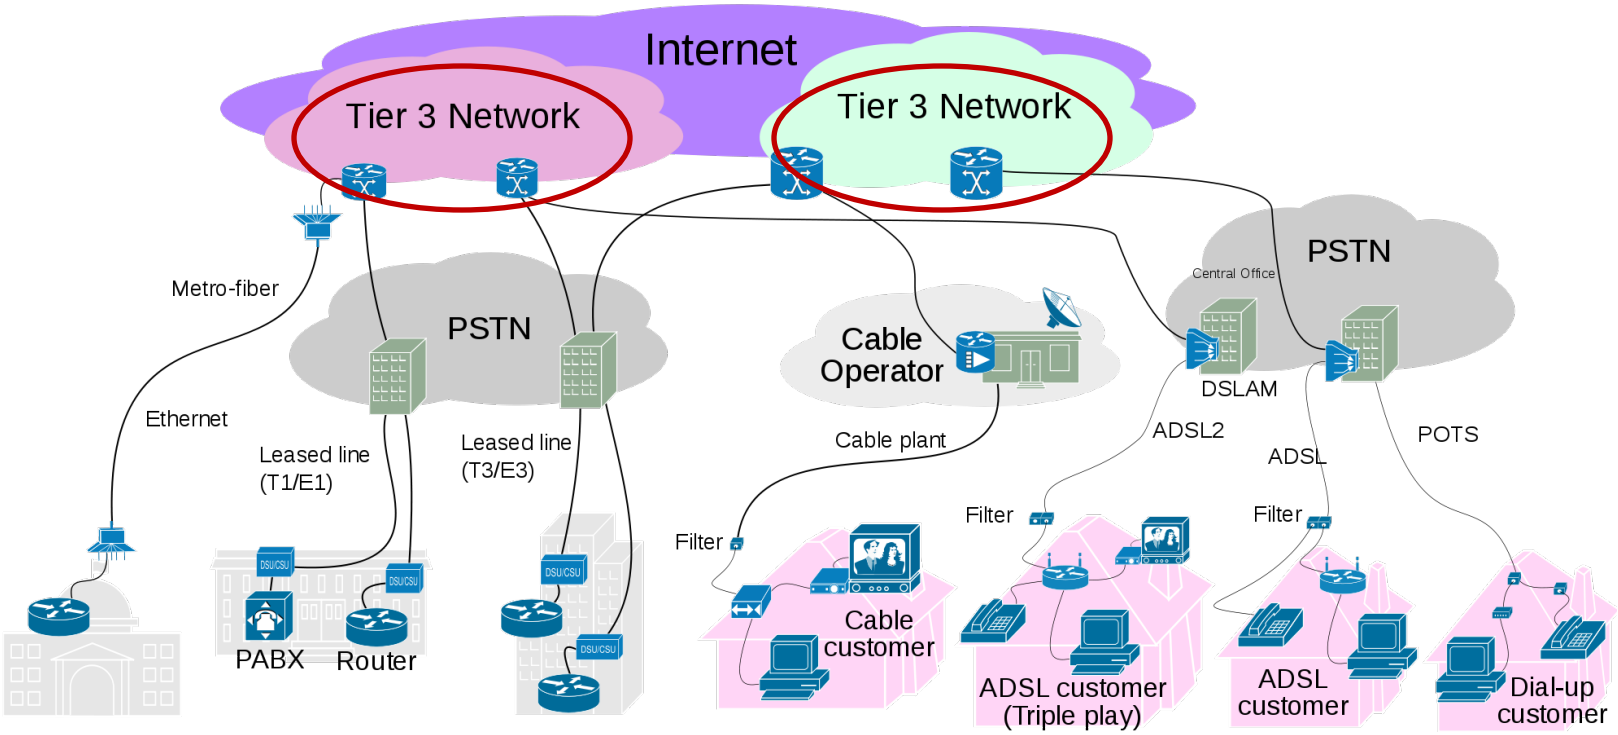

- Parallel processing has been explored
  - However, higher level analysis requiring context correlation needs sharing of state and limits scalability.
- Packet sampling has been explored
  - Creates accuracy problems

Image courtesy Wikipedia

Solutions?

- Detect in the network core (e.g. ISP)

- Detect at the network edge (e.g. home network)

Detect at Home?

- Raises a couple of questions:
  - Is there any accuracy advantage offered by home or small-office networks?
  - How do we solve the problem of management ?

Detect at Home?

- Raises a couple of questions:
  - Is there any accuracy advantage offered by home or small-office networks?
  - How do we solve the problem of management ?

Datasets for Accuracy Experiments

Benign Dataset

Dataset Type	Active Hosts	Total Packets	Duration	Packets per sec	Total Connections	Connections per sec
HOME	8	1 million	21 hrs	62.36	3,422	0.21
SOHO	29	15 million	5.5 hrs	320.4	50,082	2.61
ISP	639	28 million	10 min	12,210	304,914	523

Datasets for Accuracy Experiments

Benign Dataset

Dataset Type	Active Hosts	Total Packets	Duration	Packets per sec	Total Connections	Connections per sec
HOME	8	1 million	21 hrs	62.36	3,422	0.21
SOHO	29	15 million	5.5 hrs	320.4	50,082	2.61
ISP	639	28 million	10 min	12,210	304,914	523

Attack Dataset

Attacks	Infected Hosts	Attack Rates pkts / sec
TCP Portscan, TCP SYN Flood, UDP Flood	Around 20% of active hosts in each dataset	0.1, 1, 10, 100, 1000

Anomaly Detectors

Threshold Random Walk
IEEE S&P '04

- Uses Sequential Hypothesis Testing for detecting "horizontal" TCP portscans
- Operates on per-host basis

Anomaly Detectors

Threshold Random Walk
IEEE S&P '04

- Uses Sequential Hypothesis Testing for detecting "horizontal" TCP portscans
- Operates on per-host basis

Rate Limiting
Usenix '03

- Throttles and monitors rate of new connections
- Operates on per-host basis

Anomaly Detectors

Threshold Random Walk
IEEE S&P '04

- Uses Sequential Hypothesis Testing for detecting "horizontal" TCP portscans
- Operates on per-host basis

Rate Limiting
Usenix '03

- Throttles and monitors rate of new connections
- Operates on per-host basis

Maximum Entropy
IMC '06

- Information Theoretic Detector
- General Purpose
- Operates on time-windowed packet statistics

Anomaly Detectors

Threshold Random Walk
IEEE S&P '04

- Uses Sequential Hypothesis Testing for detecting "horizontal" TCP portscans
- Operates on per-host basis

Rate Limiting
Usenix '03

- Throttles and monitors rate of new connections
- Operates on per-host basis

Maximum Entropy
IMC '06

- Information Theoretic Detector
- General Purpose
- Operates on time-windowed packet statistics

NetAD
SAC '03

- Detects anomalous values in packet bytes
- General Purpose
- Operates on per-packet basis

TCP Portscan Results

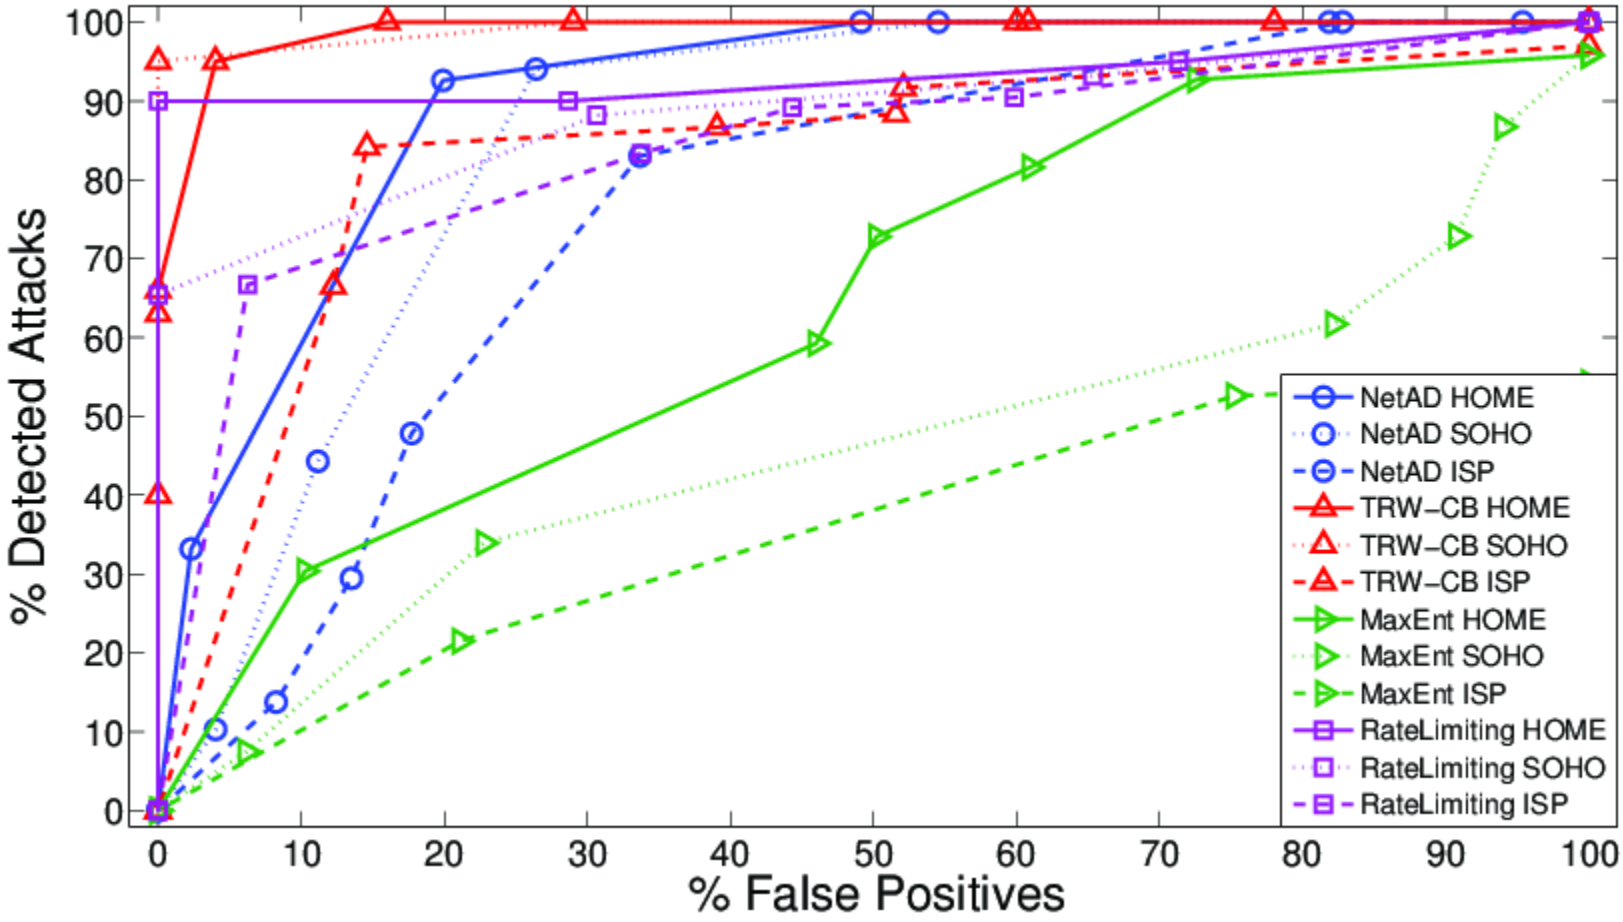

TCP Flood Results

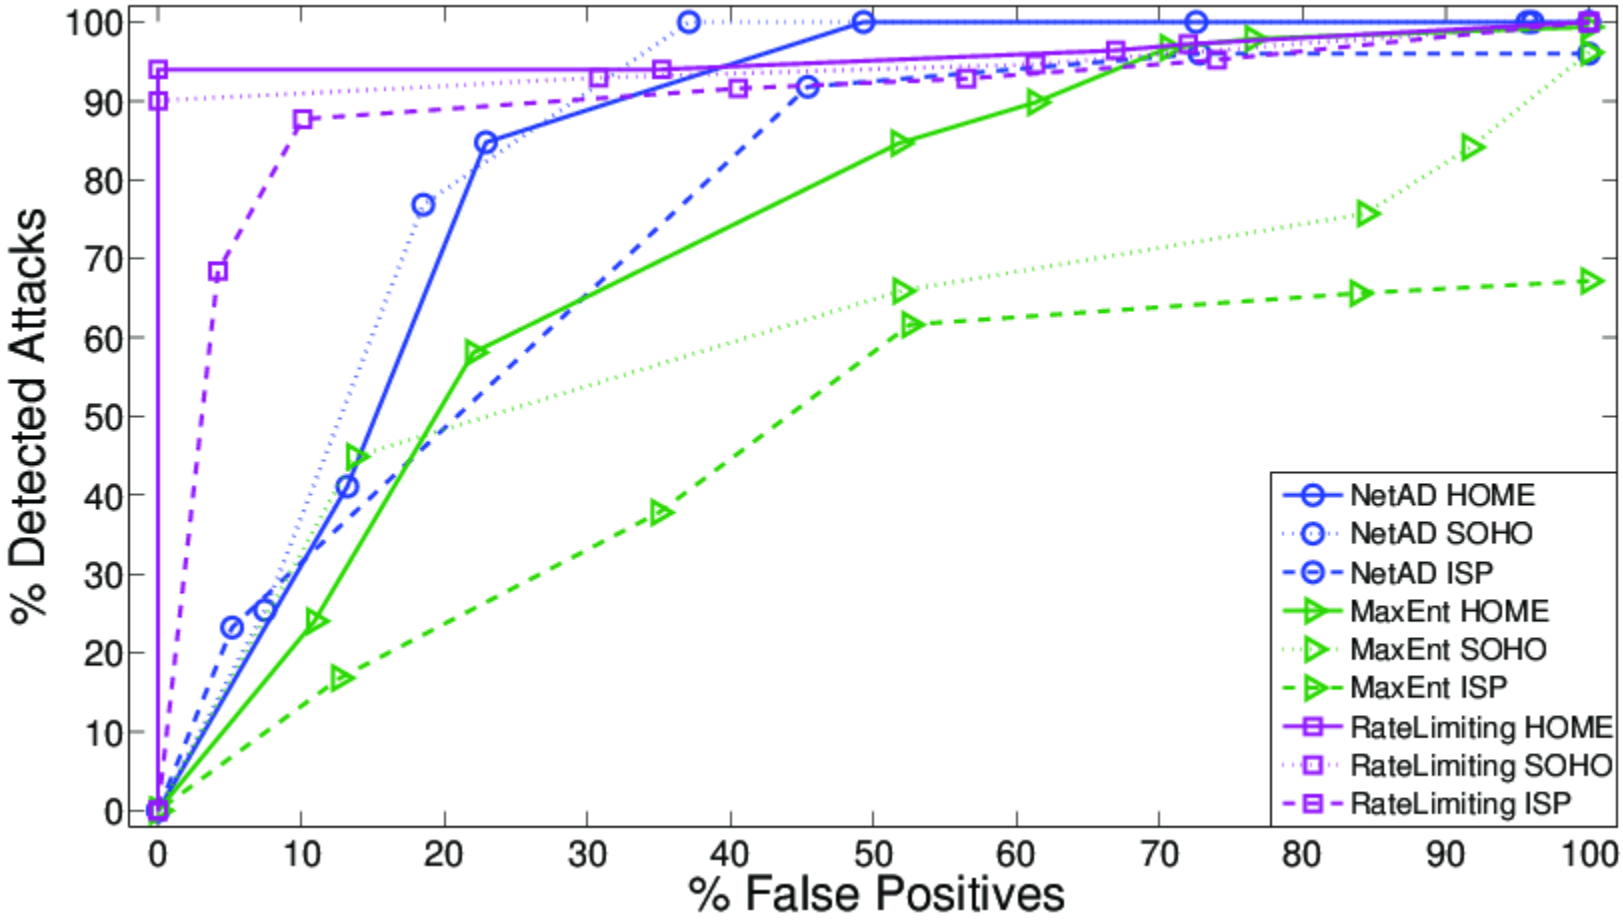

UDP Flood Results

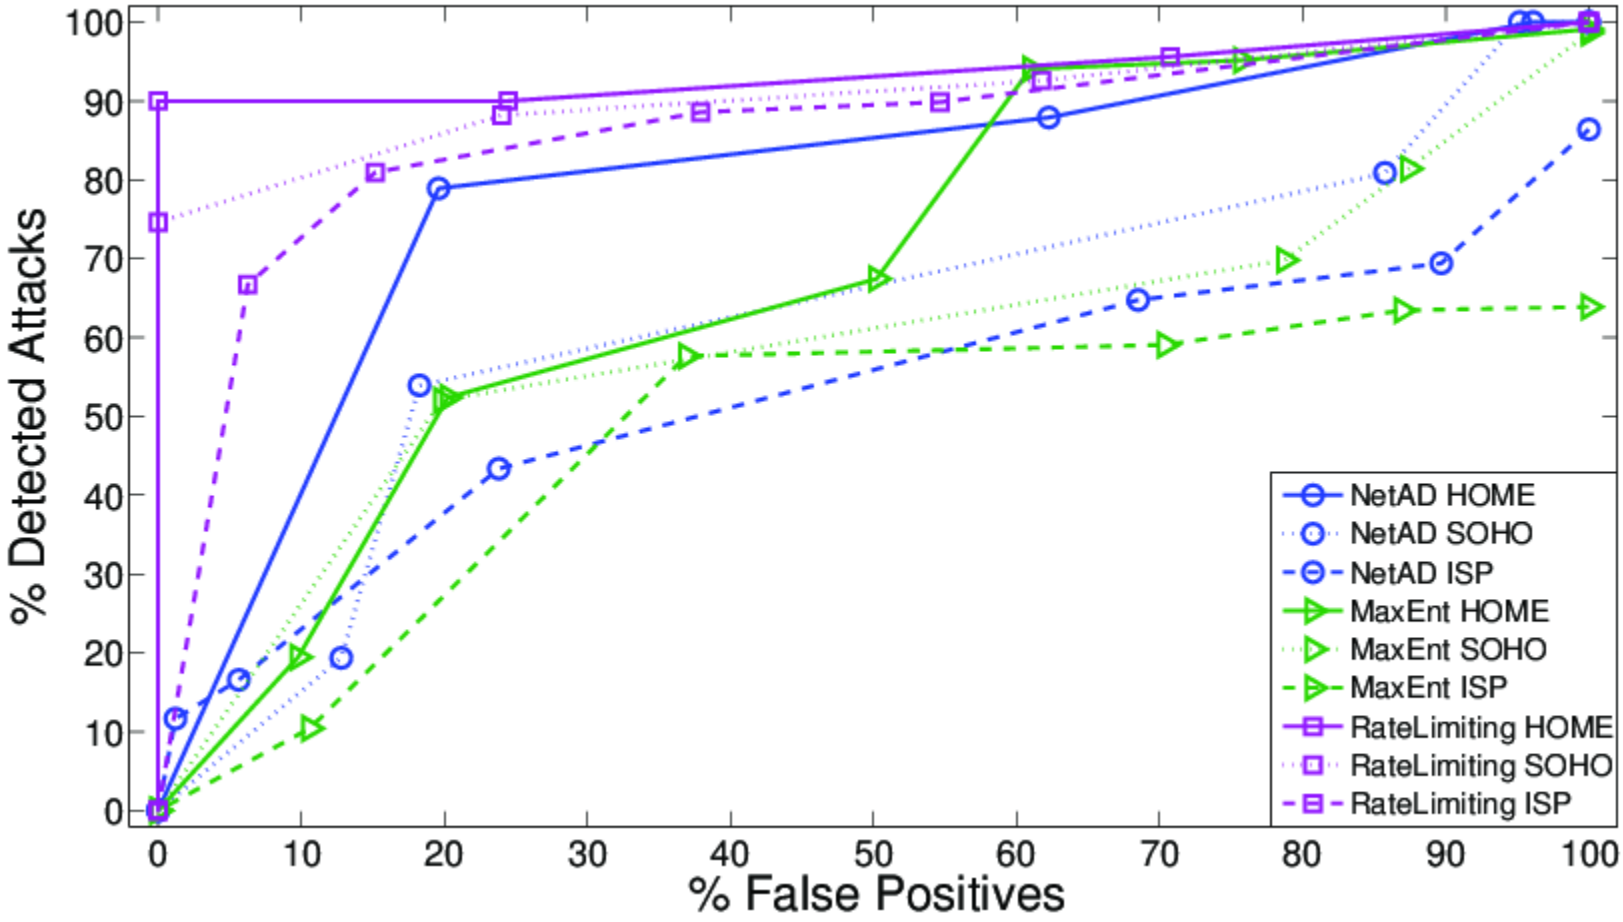

Why do Home Networks perform better?

- Less background traffic allows attacks to stand-out.
- Network Address Translation (NAT) obfuscates the ISP's perspective.
- It is possible to model "the normal" more accurately.
  - Therefore more accurate to detect genuine deviations from it.

Detect at Home?

- Raises a couple of questions:
  - What is the accuracy advantage offered by home or small-office networks (if any) ?
  - How do we solve the problem of management ?

Motivating Architecture

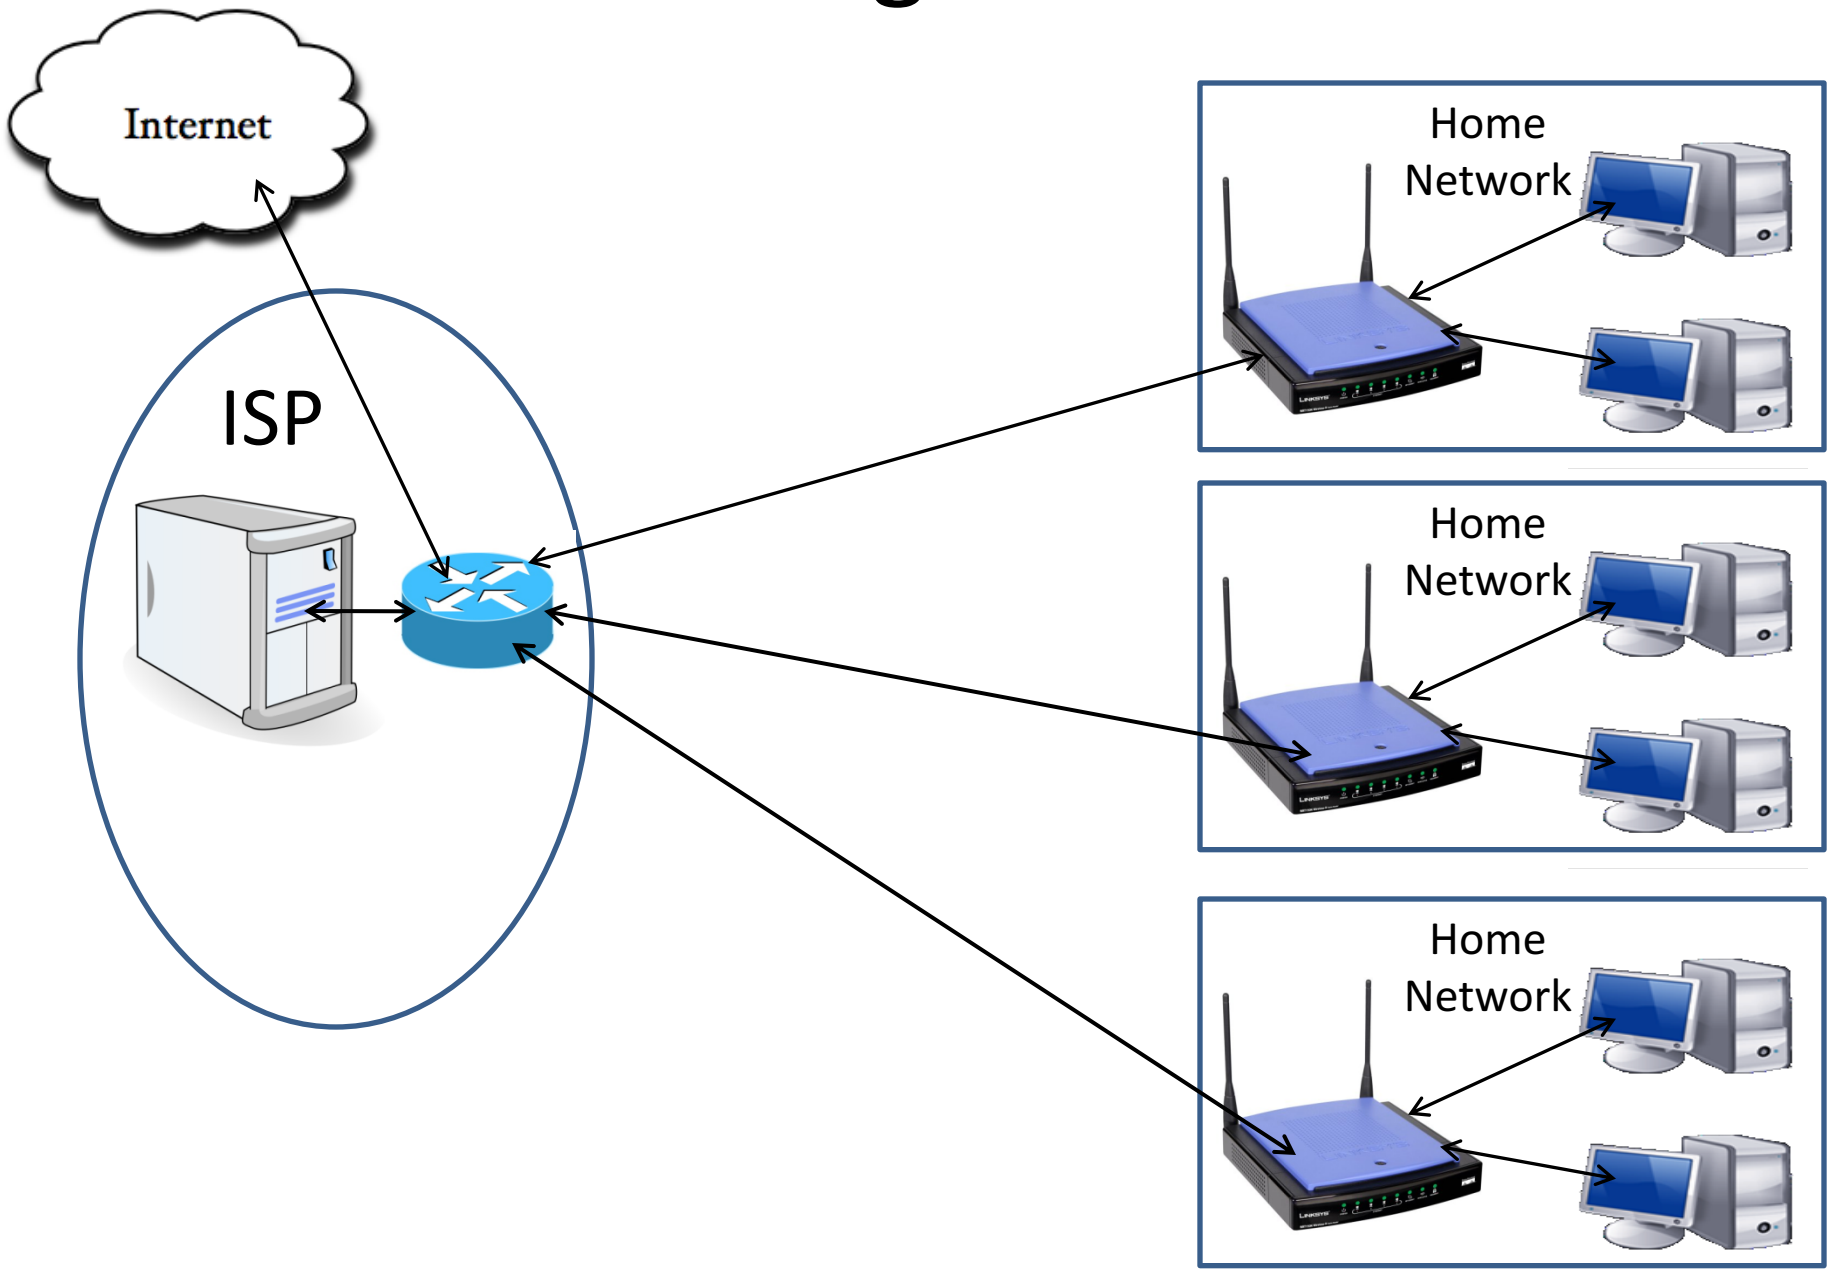

Motivating Architecture

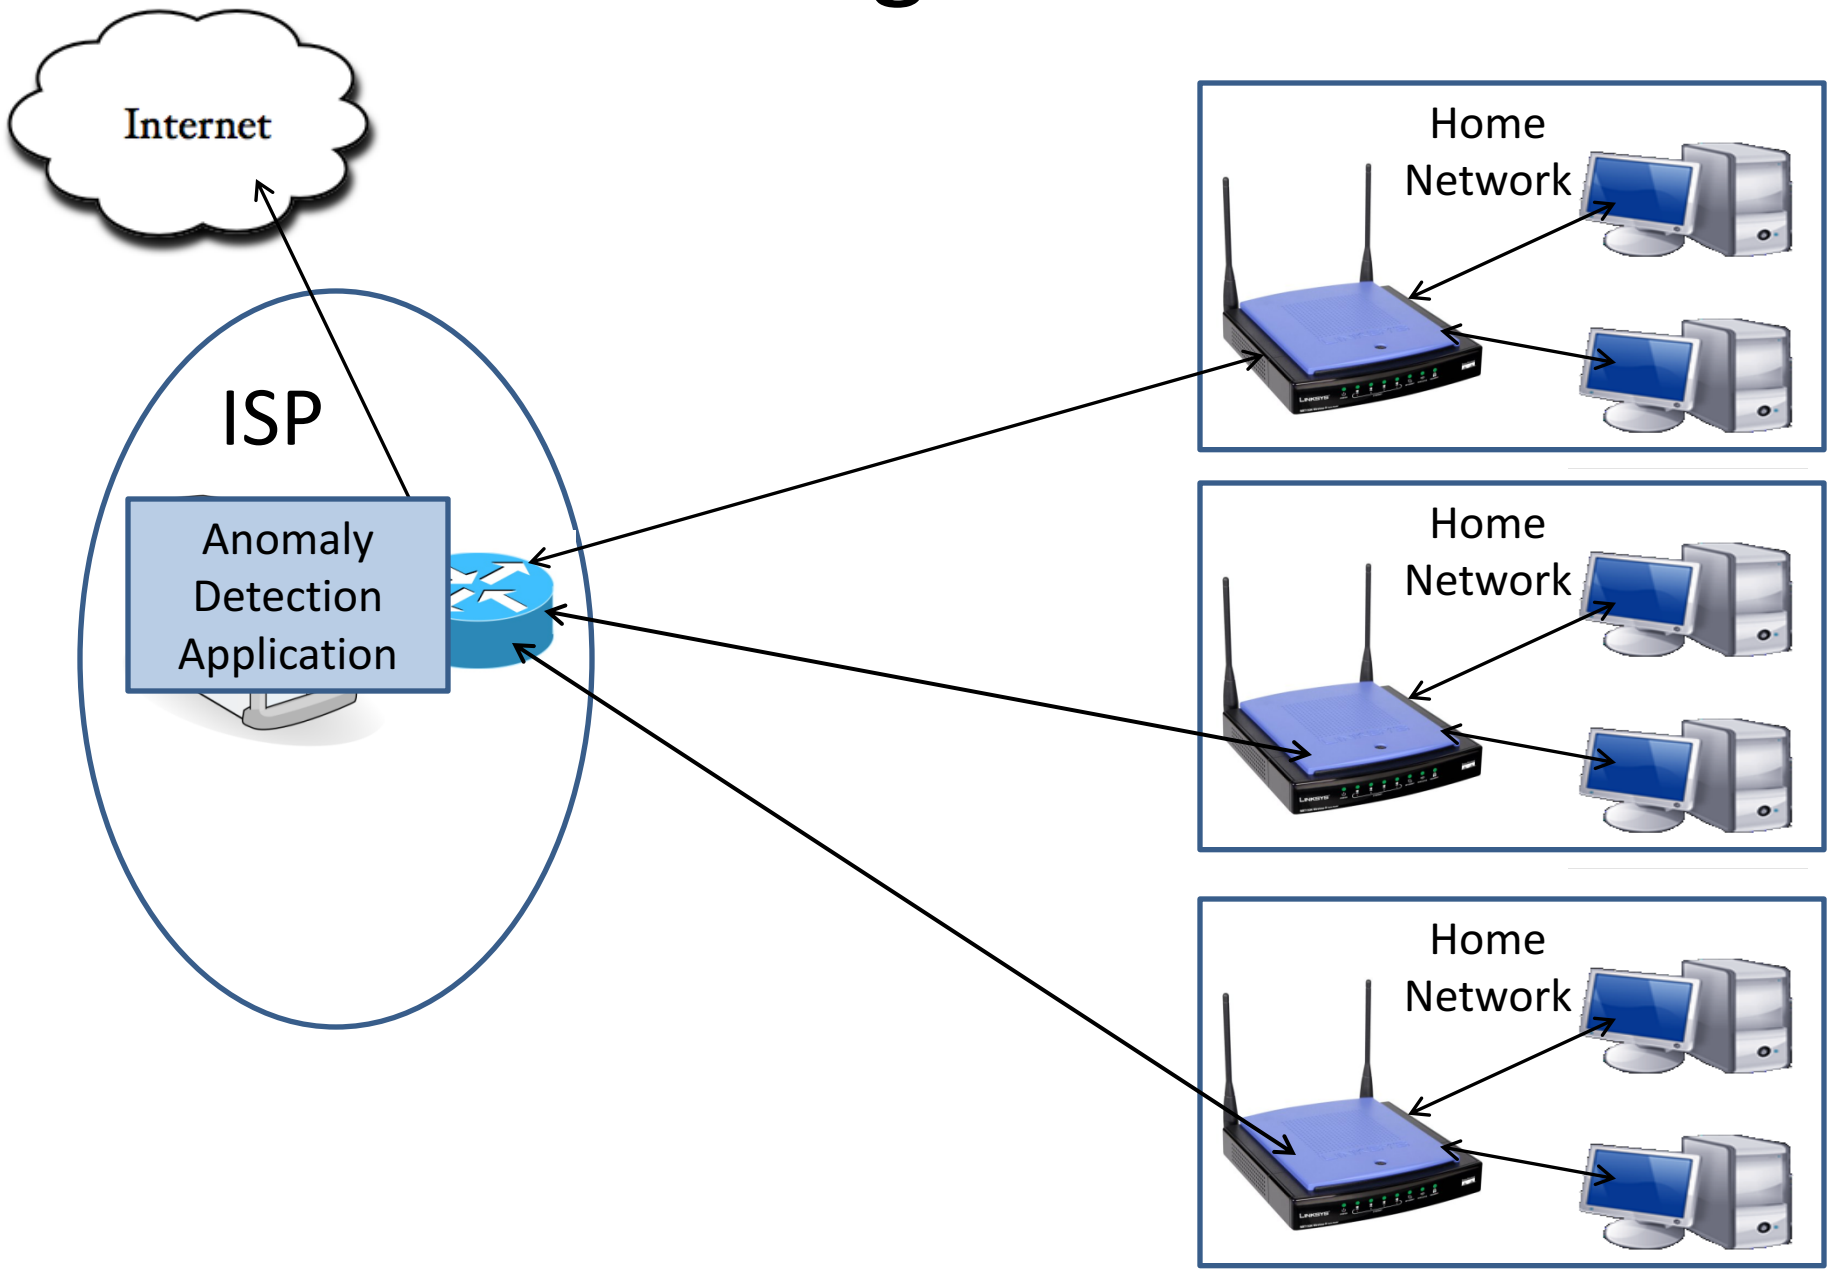

Motivating Architecture

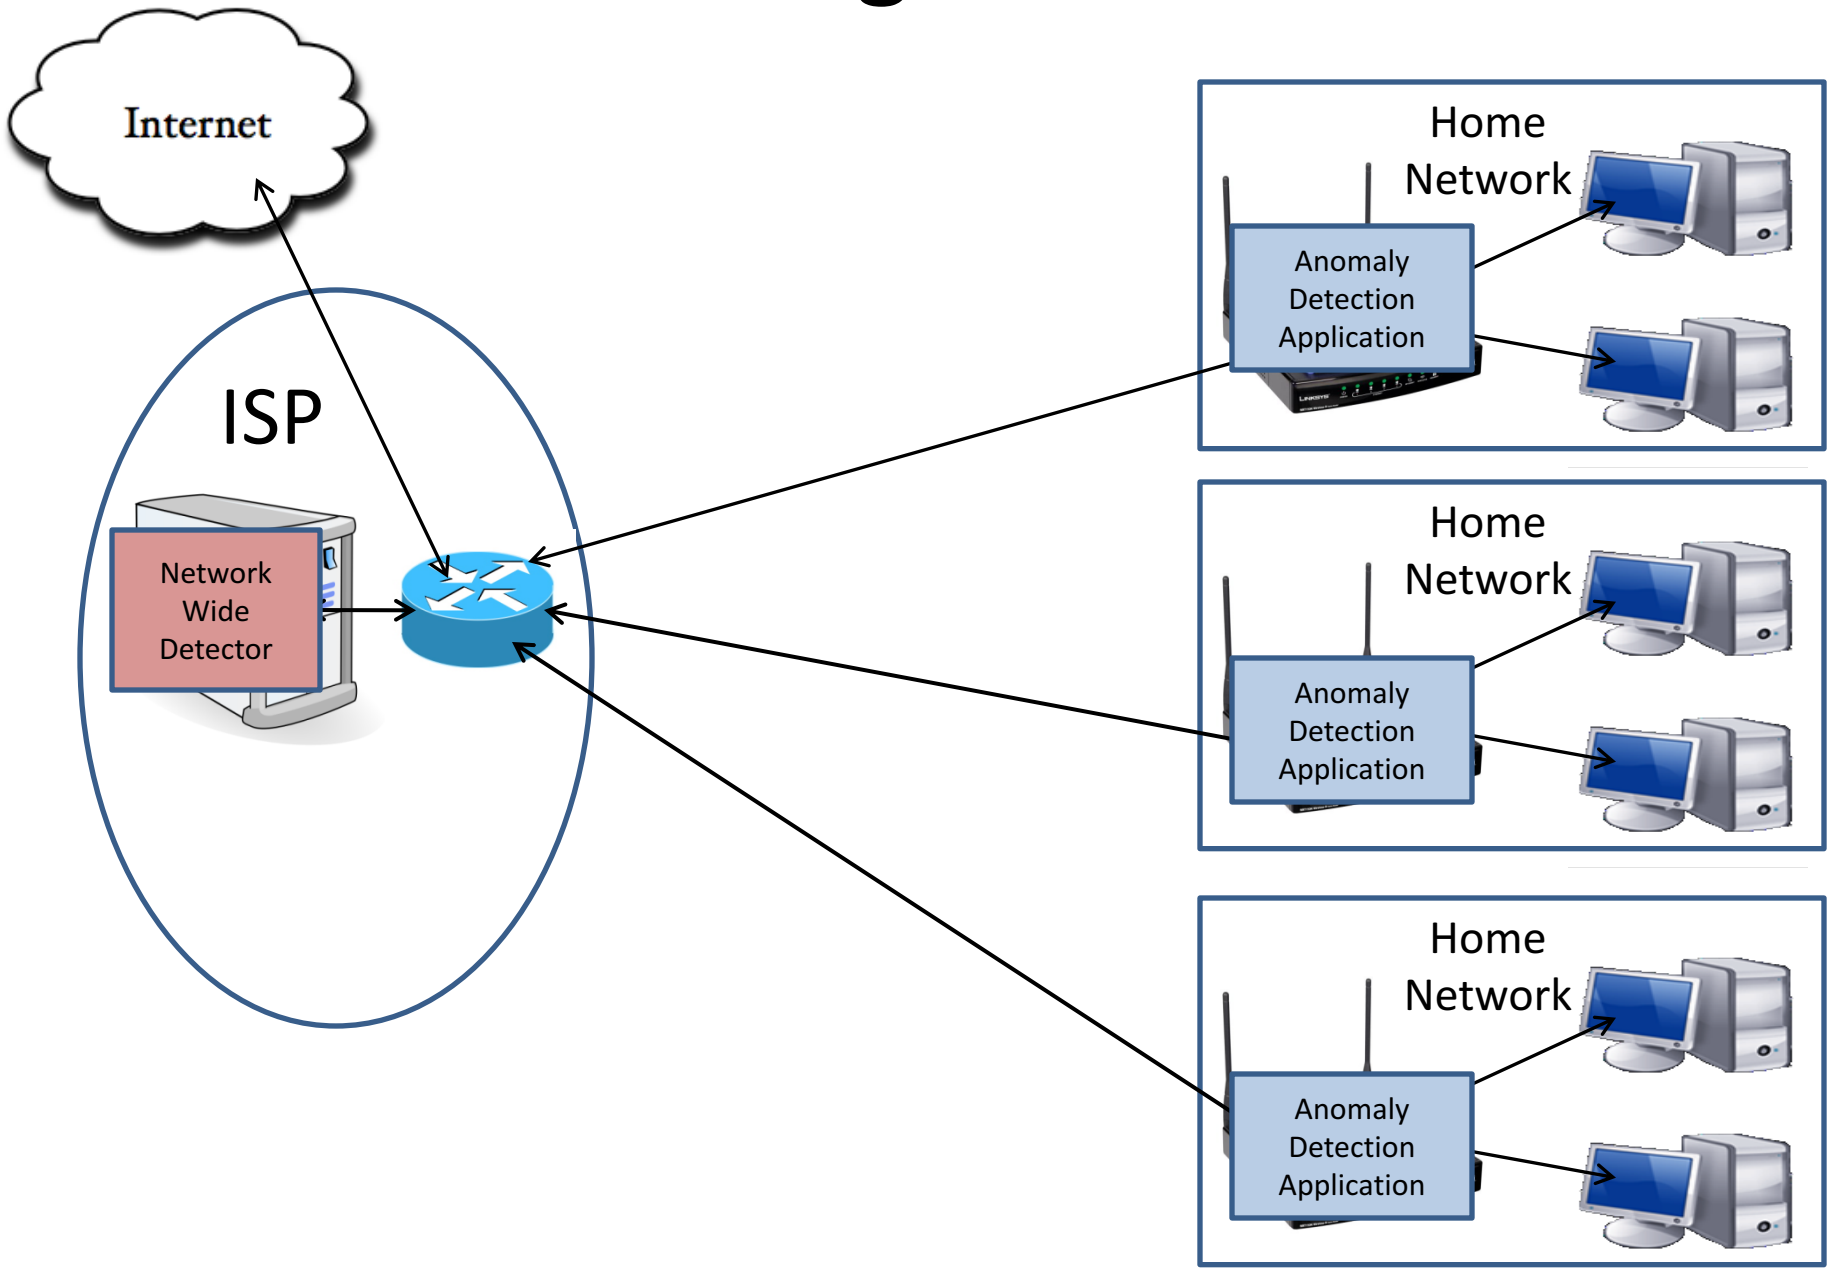

Motivating Architecture

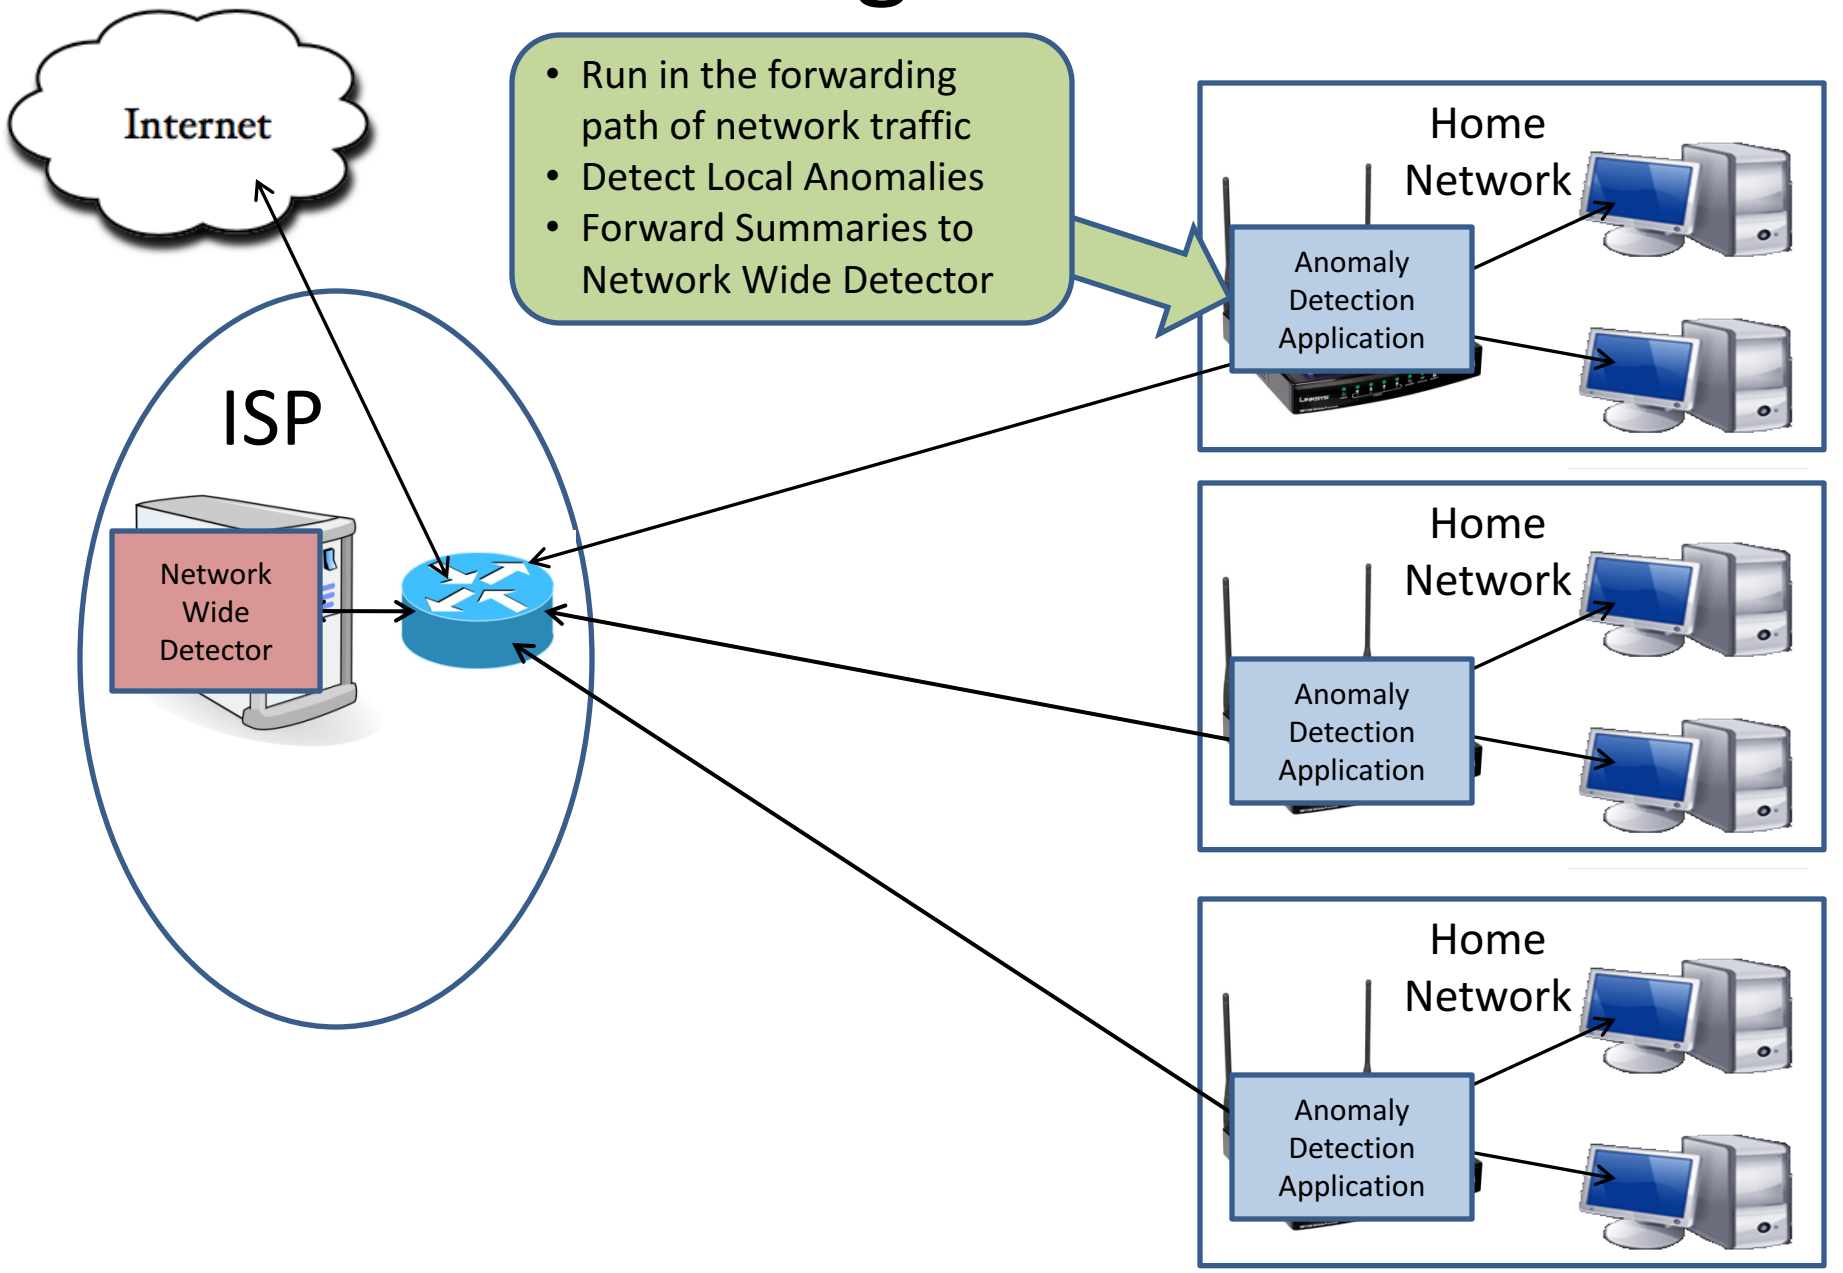

Motivating Architecture

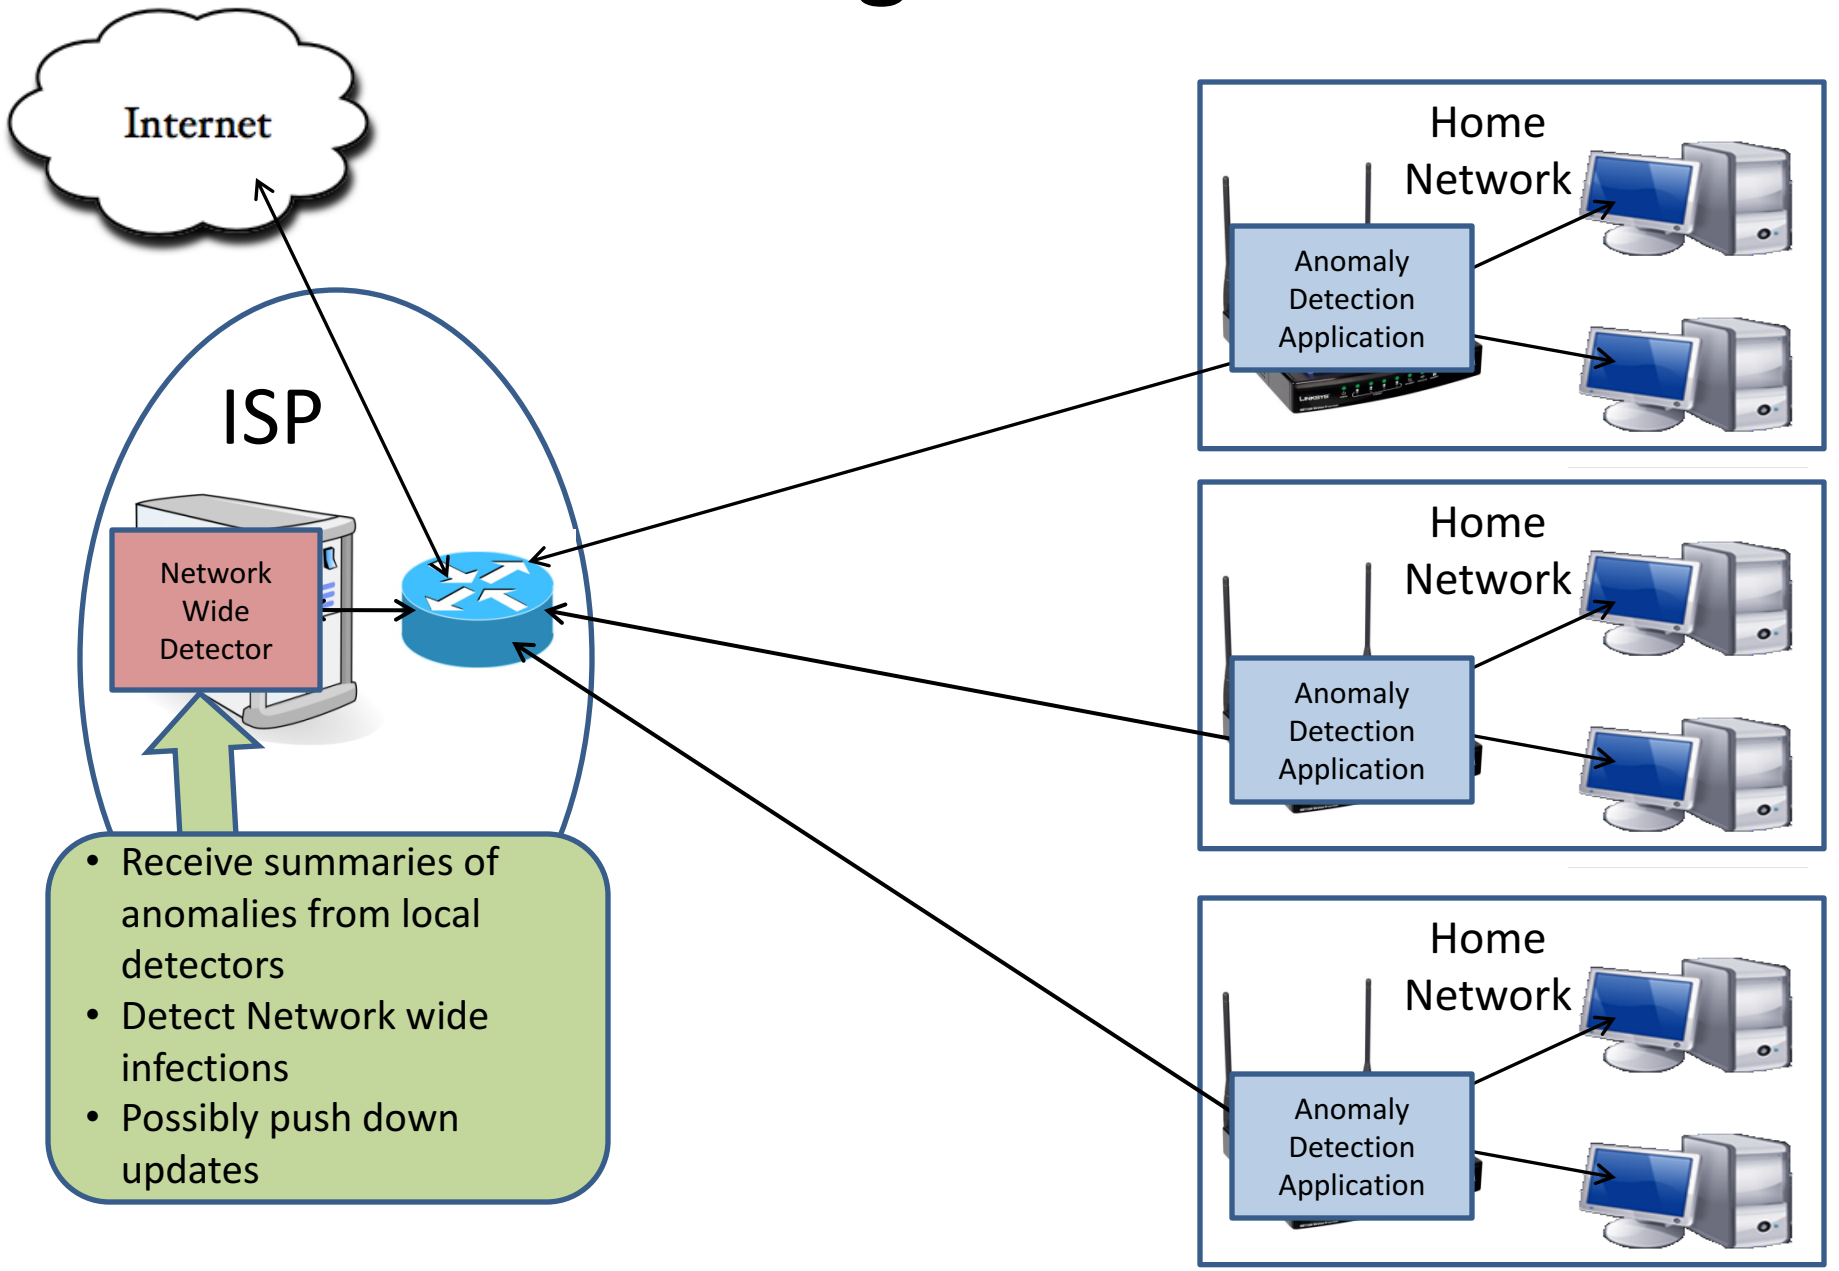

Management of security at home

- Current home routers do not provide the capability to implement such an architecture.

- Enter Software Defined Networking.

Software Defined Networking

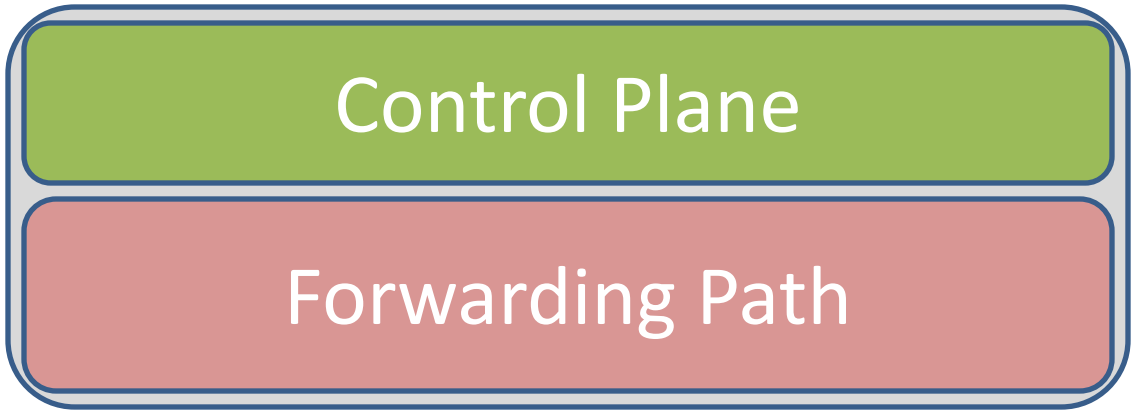

Software Defined Networking

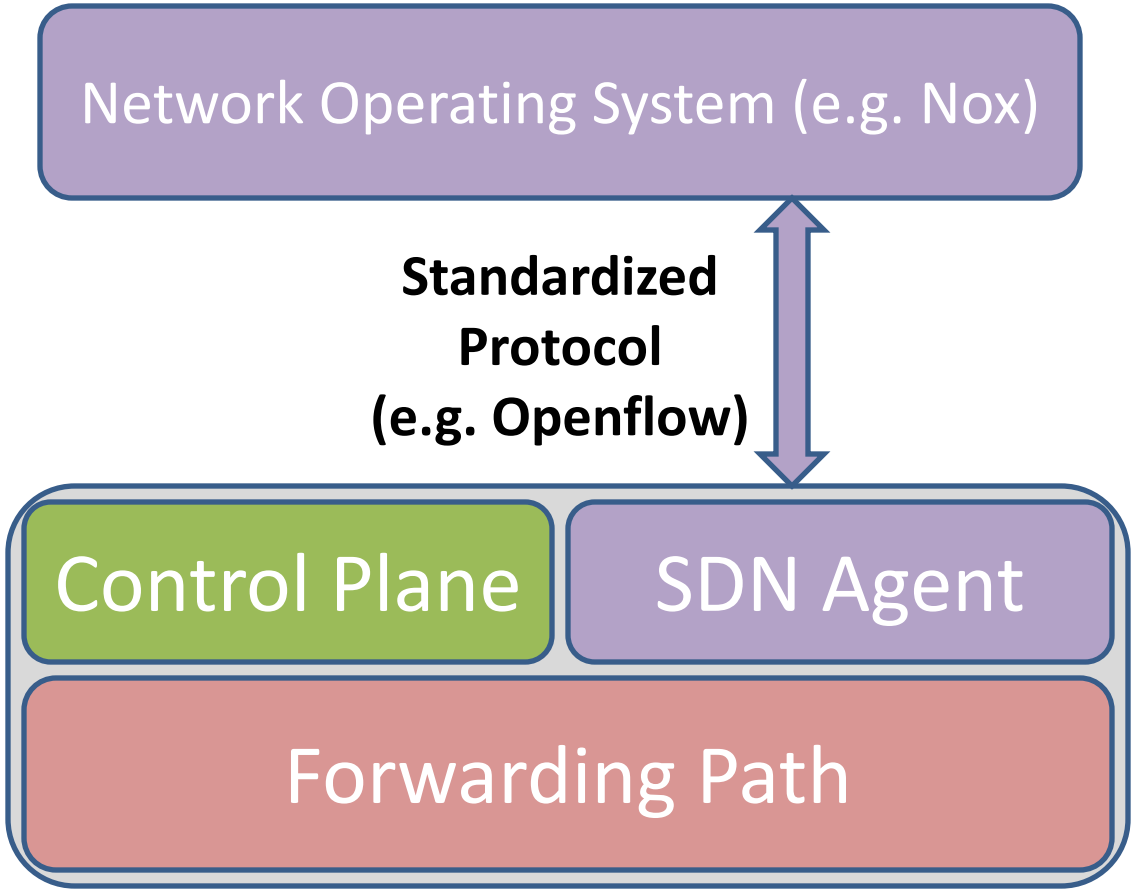

Motivating Architecture

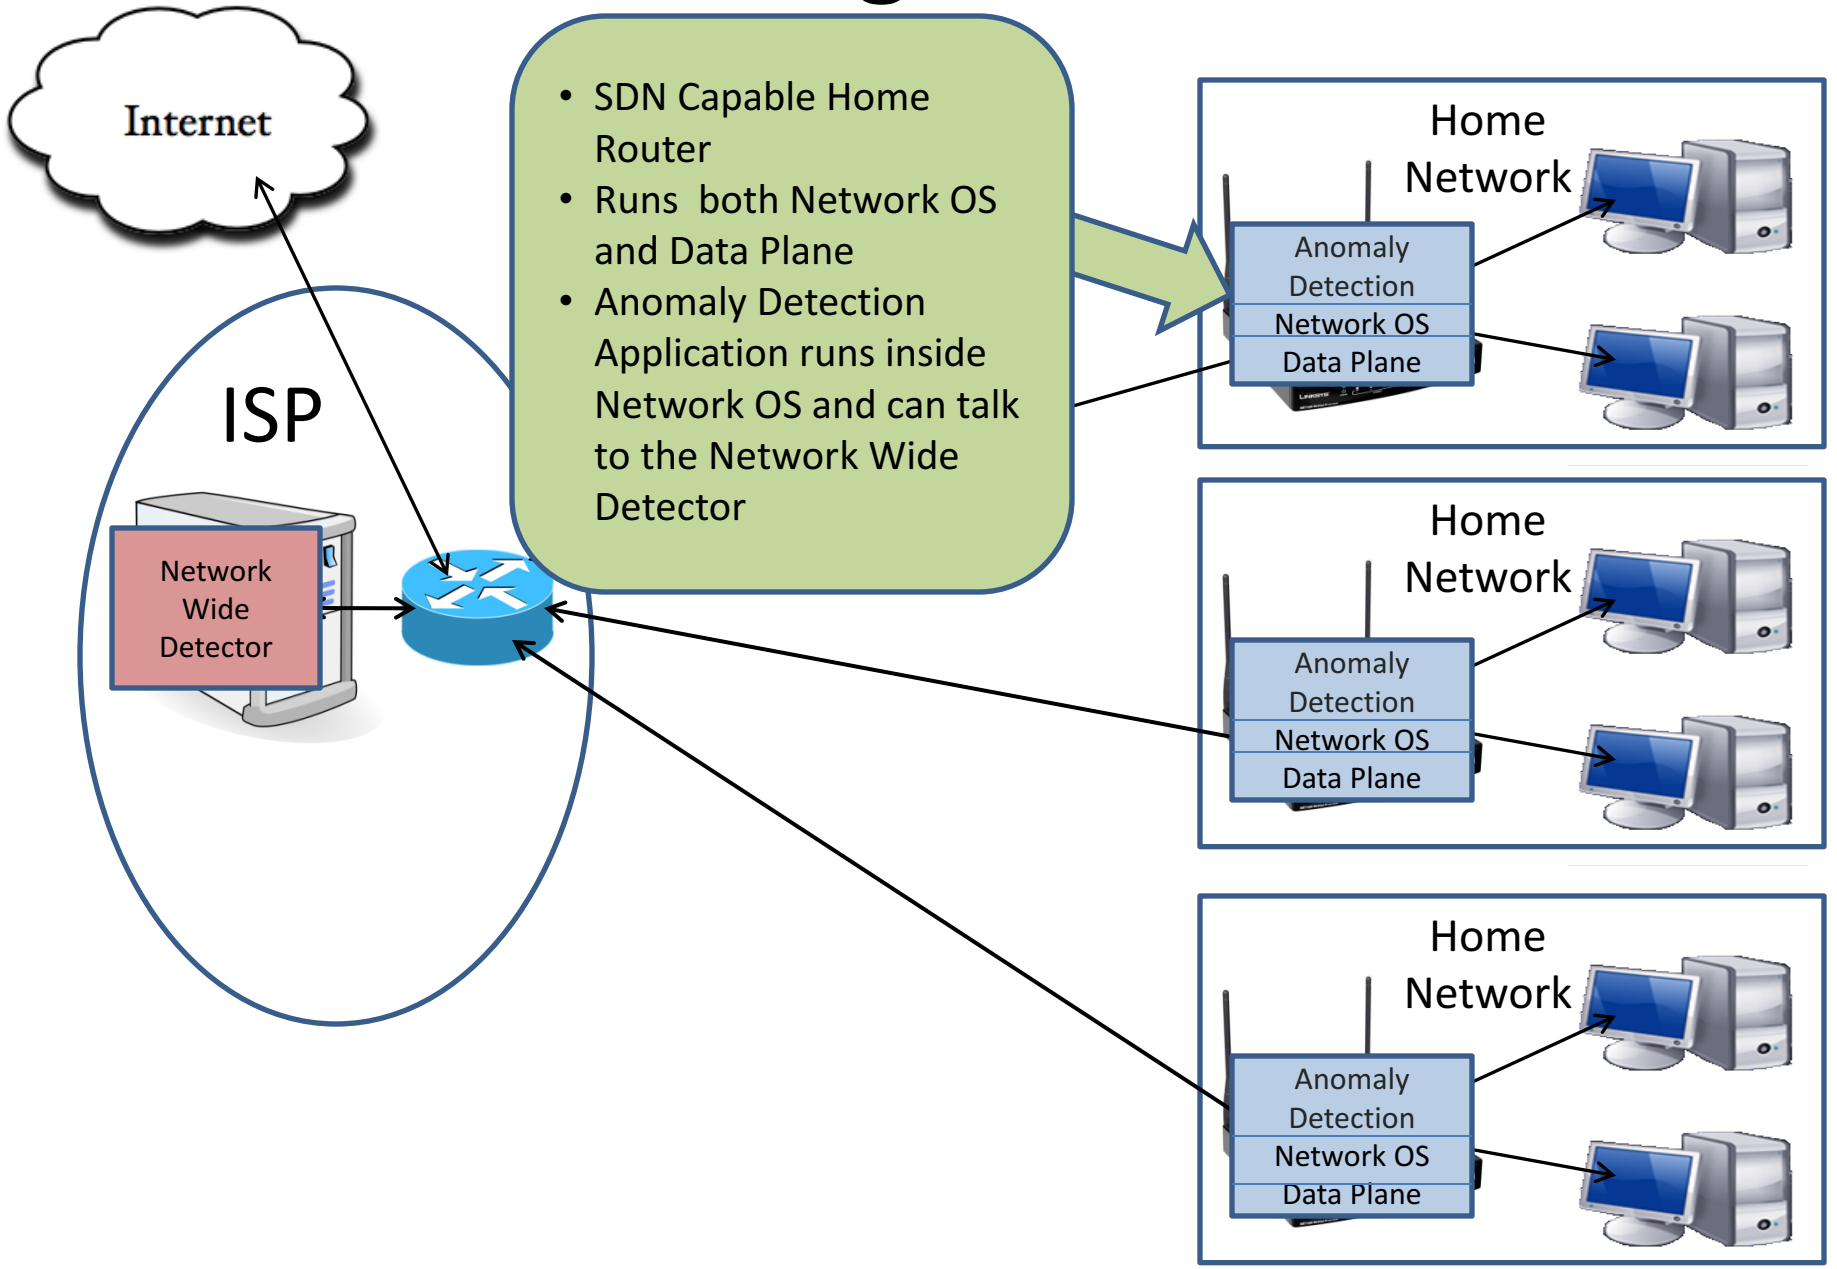

Our current scope

- Evaluate the efficiency of running an Anomaly Detection application on a home network router.
  - On top of a Network OS and Data Plane
- Use Openflow as the protocol between Network OS and Data Plane.

OpenFlow Switching

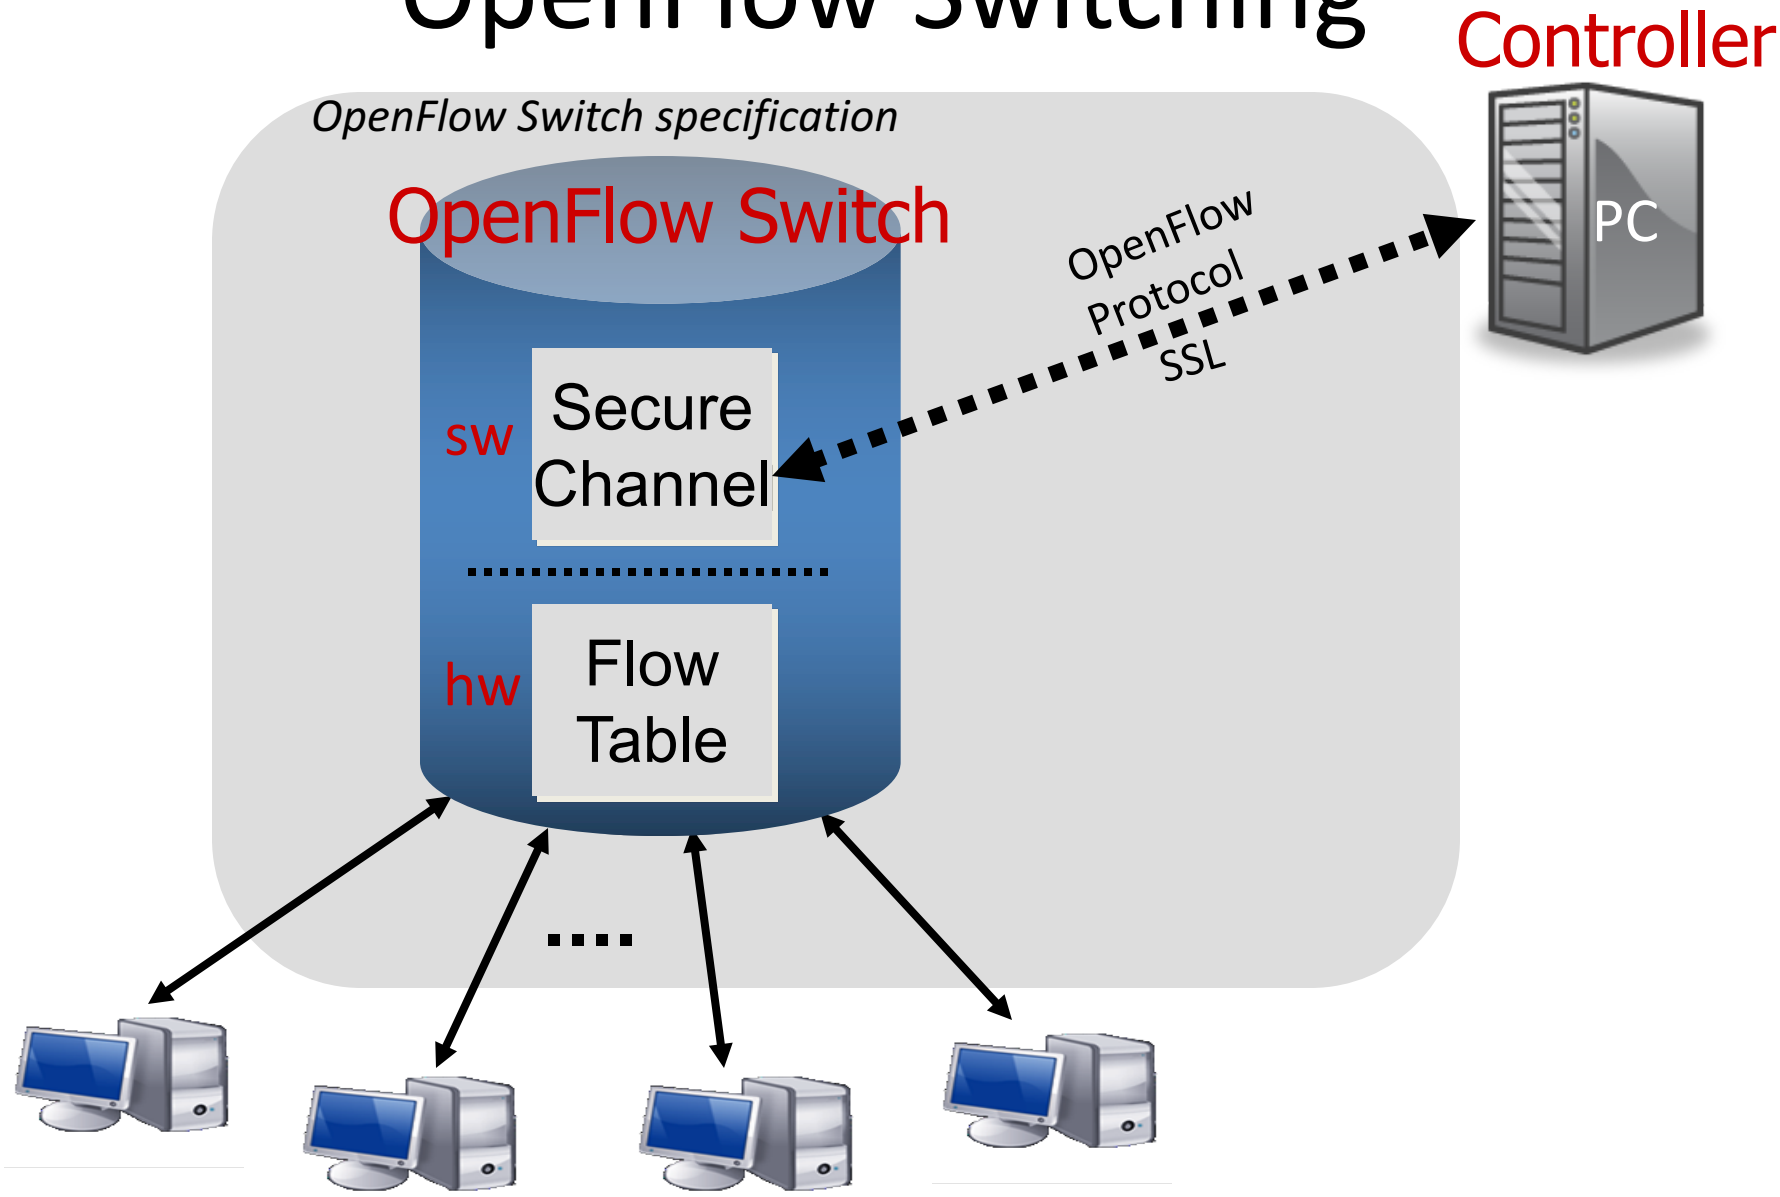

Slide Credit: Prof. Nick McKeown, Stanford

Flow Table Structure

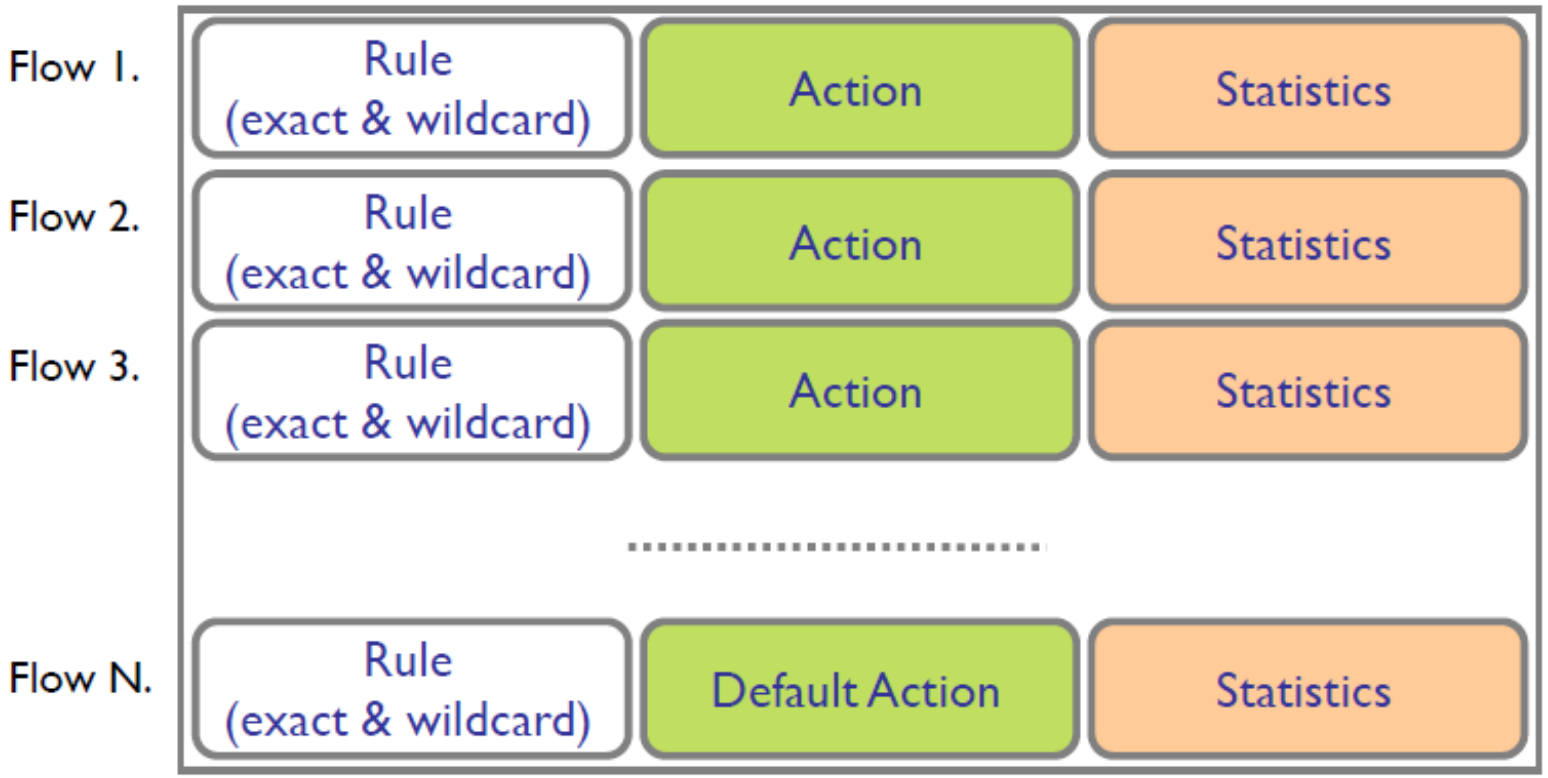

Slide Credit: Prof. Nick McKeown, Stanford

Flow Table Entry

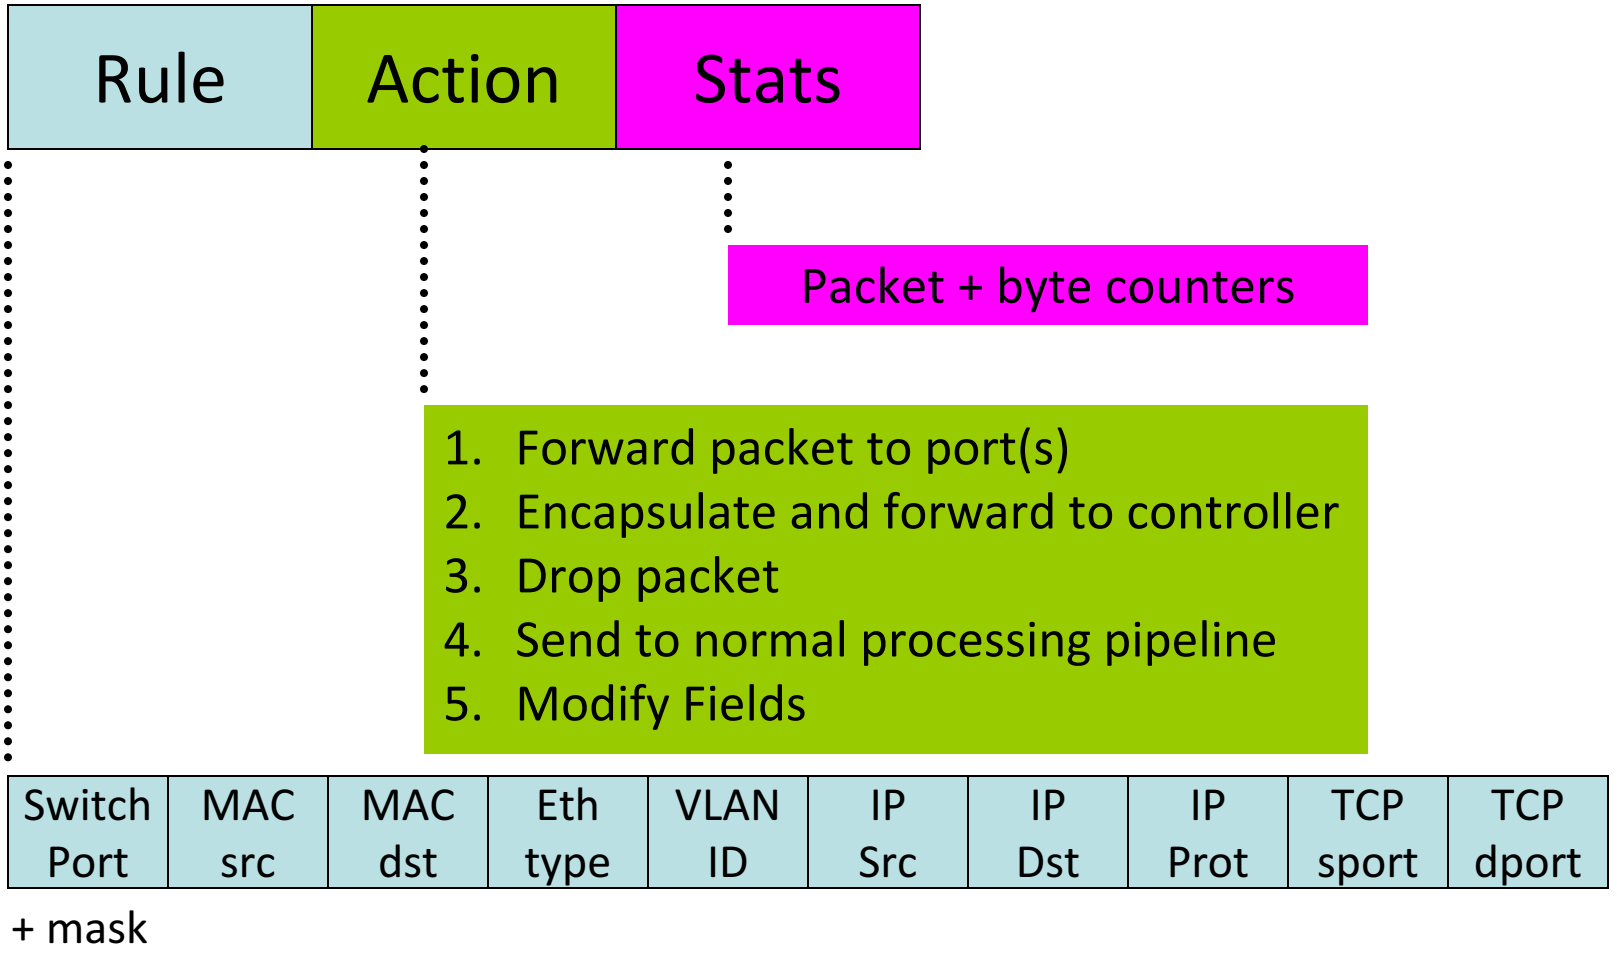

Slide Credit: Prof. Nick McKeown, Stanford

Adapting Anomaly Detection to Openflow

- What feature of Openflow allows us to efficiently adapt anomaly detection algorithms?

Observe *interesting traffic* at the anomaly detection application and let the rest pass through the fast datapath.

Observe *interesting traffic*

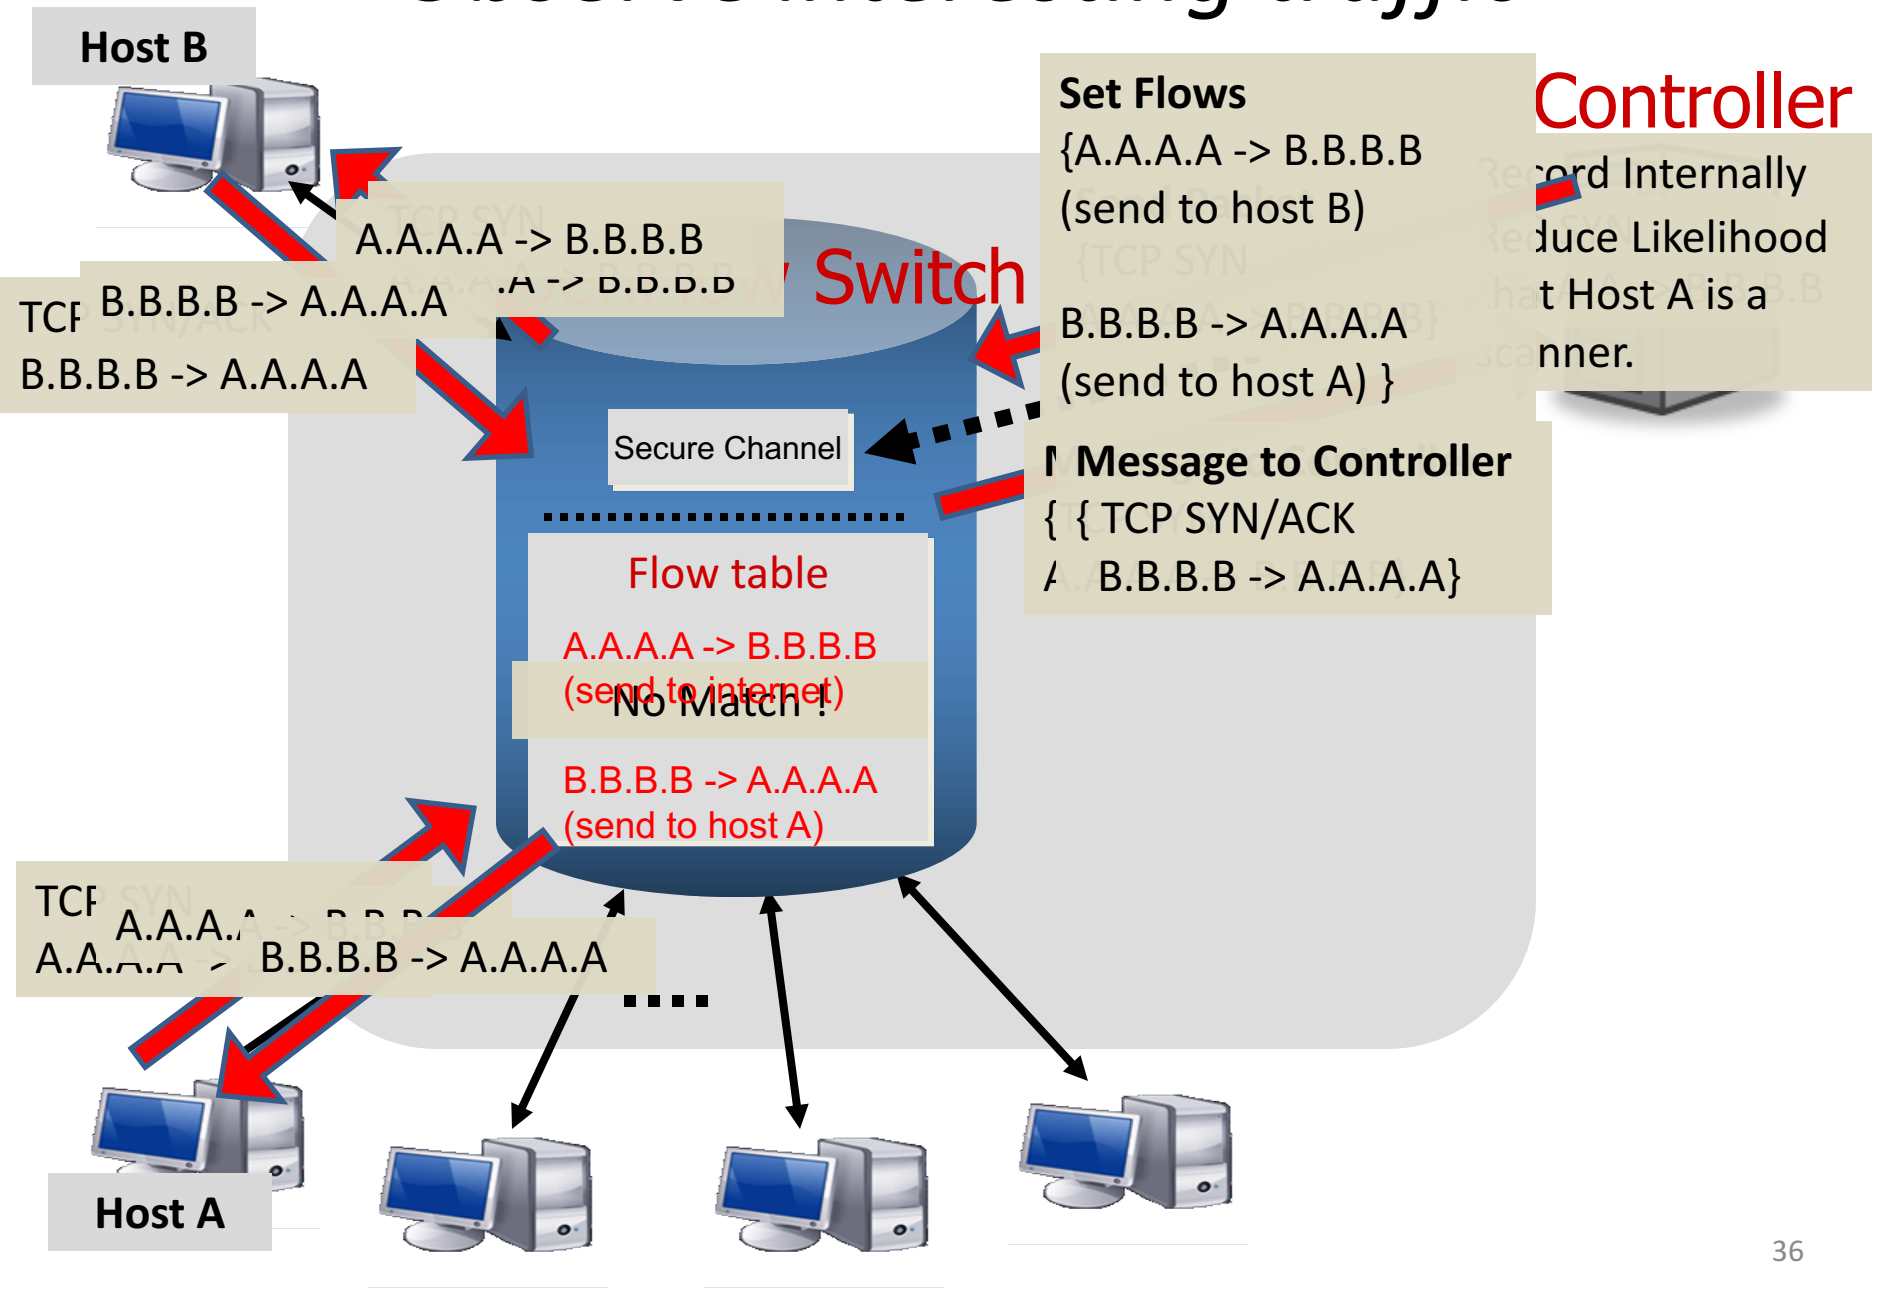

Efficiency Evaluation

	Algorithm	% of total packets at Controller
HOME	TRW	1.15 %
	Rate Limit	1.00 %
	Max Ent	2.48 %
	NetAD	3.46 %

Efficiency Evaluation

	Algorithm	% of total packets at Controller	Pkt rate at Controller per sec
HOME	TRW	1.15 %	0.73
	Rate Limit	1.00 %	0.64
	Max Ent	2.48 %	1.58
	NetAD	3.46 %	2.21

Efficiency Evaluation

	Algorithm	% of total packets at Controller	Pkt rate at Controller per sec	Avg. entries in Flow Table
HOME	TRW	1.15 %	0.73	16.11
	Rate Limit	1.00 %	0.64	16.69
	Max Ent	2.48 %	1.58	39.72
	NetAD	3.46 %	2.21	24.60

Efficiency Evaluation

	Algorithm	% of total packets at Controller	Pkt rate at Controller per sec	Avg. entries in Flow Table	Peak entries in Flow Table
HOME	TRW	1.15 %	0.73	16.11	70
	Rate Limit	1.00 %	0.64	16.69	59
	Max Ent	2.48 %	1.58	39.72	261
	NetAD	3.46 %	2.21	24.60	107

Efficiency Evaluation

	Algorithm	% of total packets at Controller	Pkt rate at Controller per sec	Avg. entries in Flow Table	Peak entries in Flow Table
HOME	TRW	1.15 %	0.73	16.11	70
	Rate Limit	1.00 %	0.64	16.69	59
	Max Ent	2.48 %	1.58	39.72	261
	NetAD	3.46 %	2.21	24.60	107
SOHO	TRW	0.37 %	2.91	42.33	71
	Rate Limit	0.56 %	4.43	38.28	64
	Max Ent	1.26 %	1.00	172.60	408
	NetAD	1.07 %	8.47	74.68	196

Nox Box

Specifications
Open vSwitch v1.0
NOX Controller
Voyage Linux
500 Mhz CPU

CPU Usage of NoxBox
Home Dataset

Data Rate	Average CPU Usage (%)			
	TRW	Rate Limit	NetAD	Max Ent
1 Mbps	1.86	2.1	2.94	3.09

CPU Usage of NoxBox
Home Dataset

Data Rate	Average CPU Usage (%)			
	TRW	Rate Limit	NetAD	Max Ent
1 Mbps	1.86	2.1	2.94	3.09
10 Mbps	6.70	8.47	10.43	18.43

CPU Usage of NoxBox
Home Dataset

Data Rate	Average CPU Usage (%)			
	TRW	Rate Limit	NetAD	Max Ent
1 Mbps	1.86	2.1	2.94	3.09
10 Mbps	6.70	8.47	10.43	18.43
50 Mbps	17.54	18.87	19.11	28.26

Summary

- The home network allows better accuracy for detection of anomalies.

- Software Defined Networking can allow the development of a solution which:
  - Pushes down some of the processing to the home network router
  - Makes it possible to have remote management of Network Security at home.

Questions?

EXTRA SLIDES

TCP Portscan High-Rate

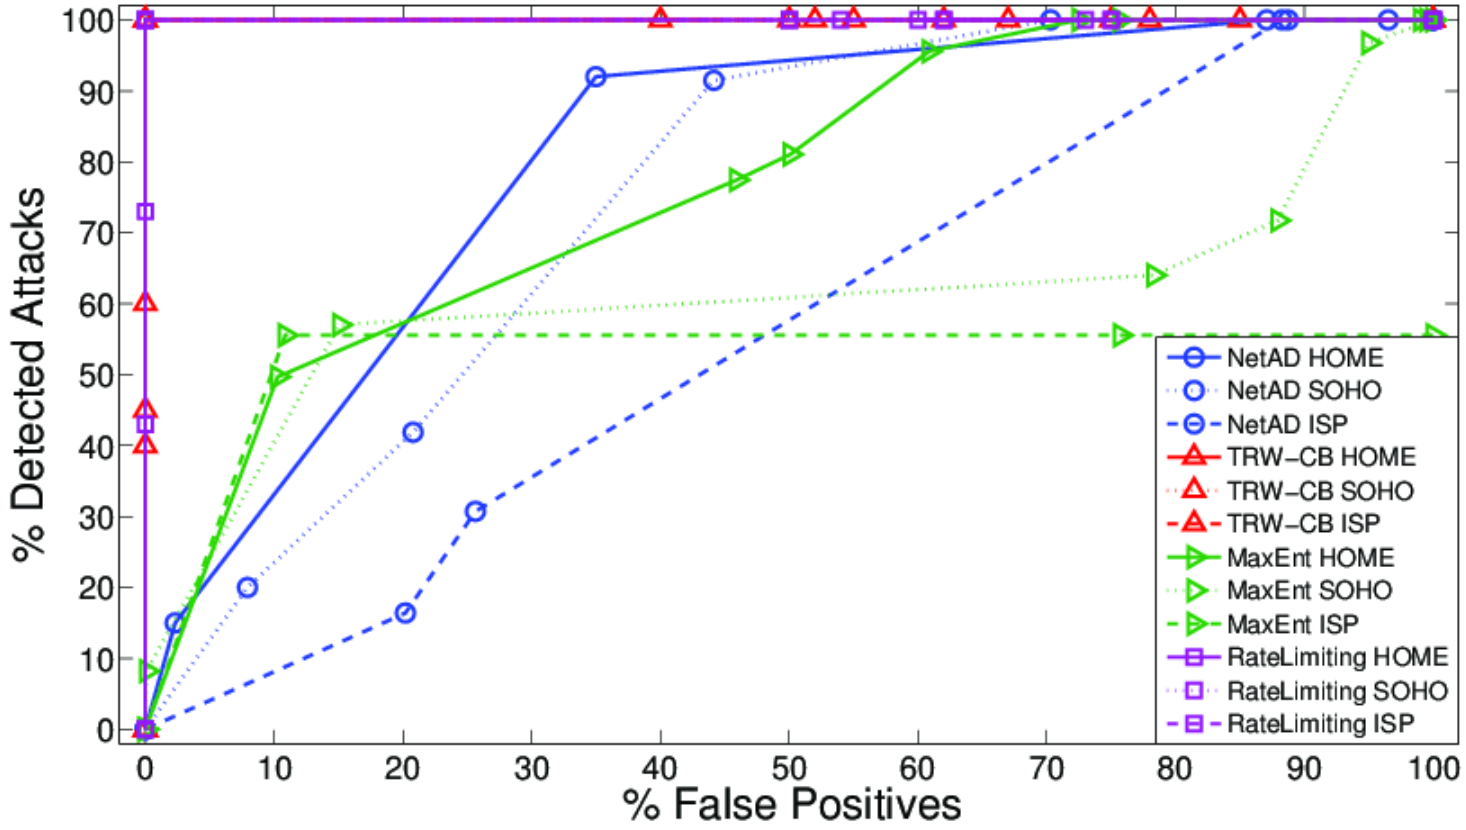

TCP Portscan Low-Rate

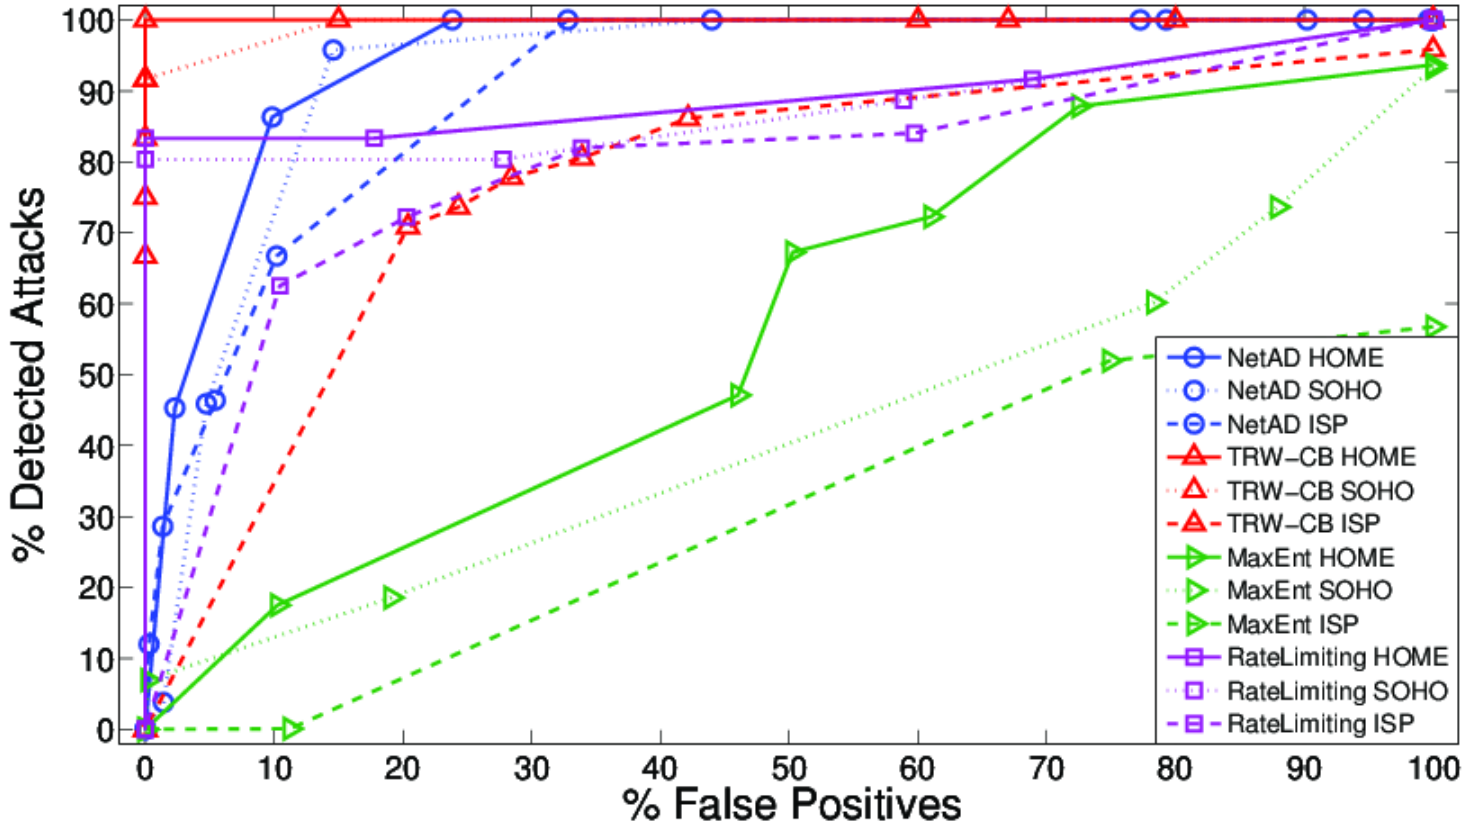

TCP Flood Low-Rate

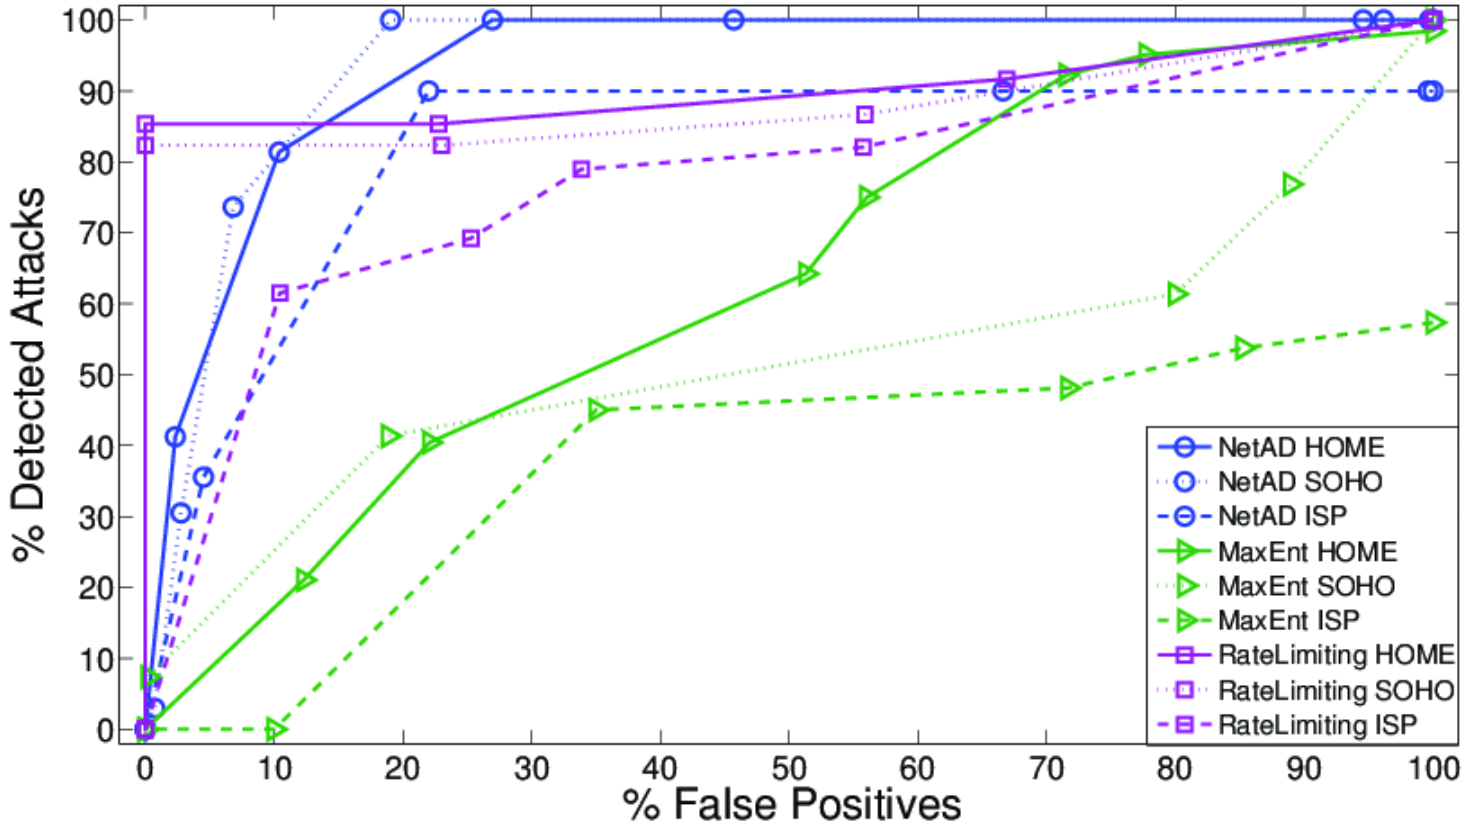

TCP Flood High-Rate

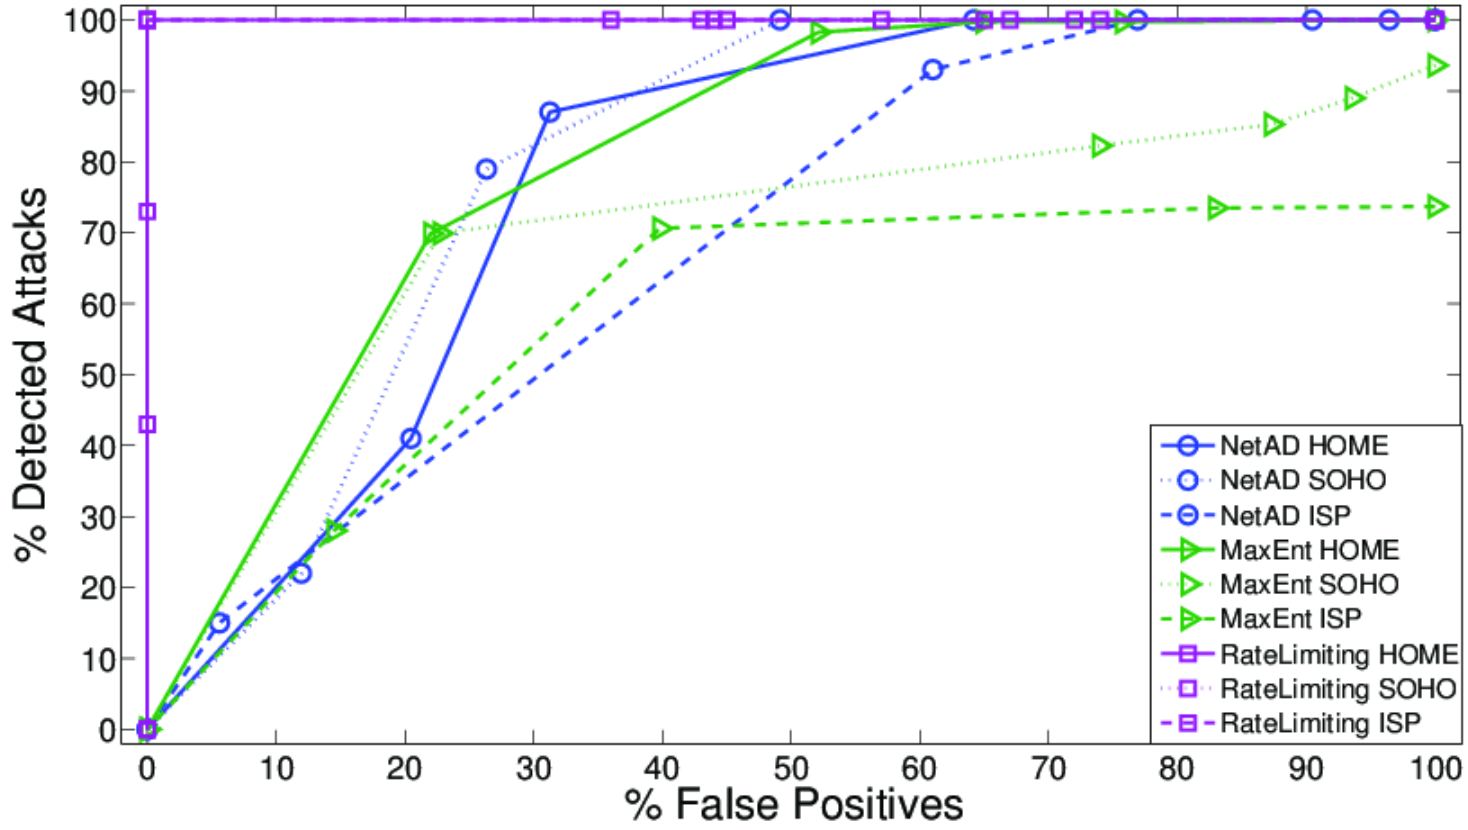

UDP Flood Low-Rate

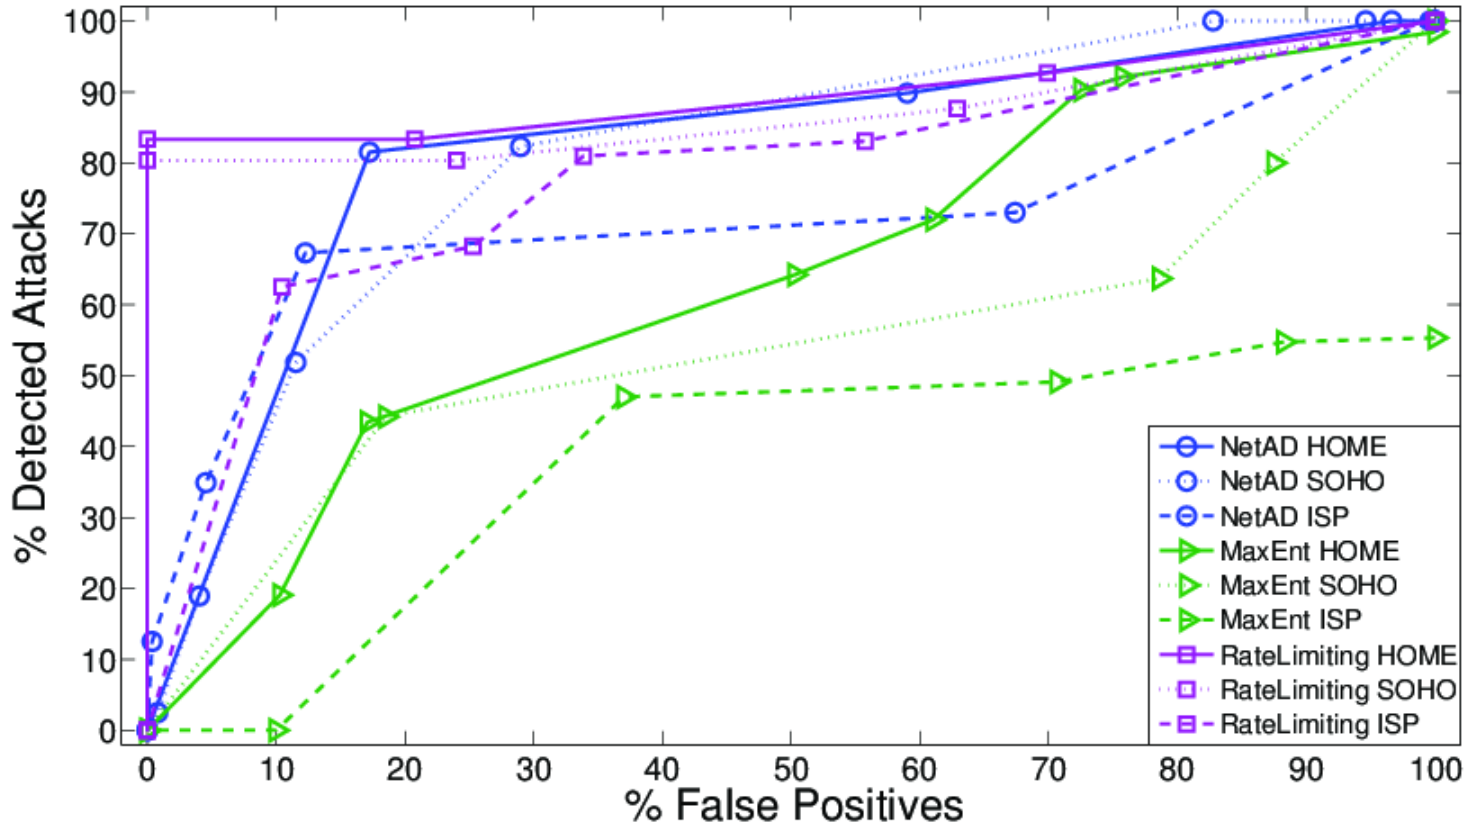

UDP Flood High-Rate