The Crossfire Attack

Min Suk Kang Soo Bum Lee Virgil D. Gligor

ECE Department and CyLab,
Carnegie Mellon University

May 20 2013

Old: DDoS Attacks against *Single Servers*

- **typical attack**: floods server with HTTP, UDP, SYN, ICMP… packets
- **persistence**

  - maximum: **2.5 days** (outlier: **81 days**)

  - average: **1.5 days**

Adversary's Challenge:

DDoS Attacks are <u>either</u> Persistent <u>or</u> Scalable to N Servers

- ***N* x** traffic to **1 server** => high-intensity traffic **triggers network detection**
- **detection not triggered** => low-intensity traffic is insufficient for ***N*** servers

Example: “Spamhaus” Attack (2013)

- Adversary: DDoS -> **1** Spamhaus ***Server***

  3/16 – 3/18: ~ 10 Gbps

  persistent: ~ **2.5 days**

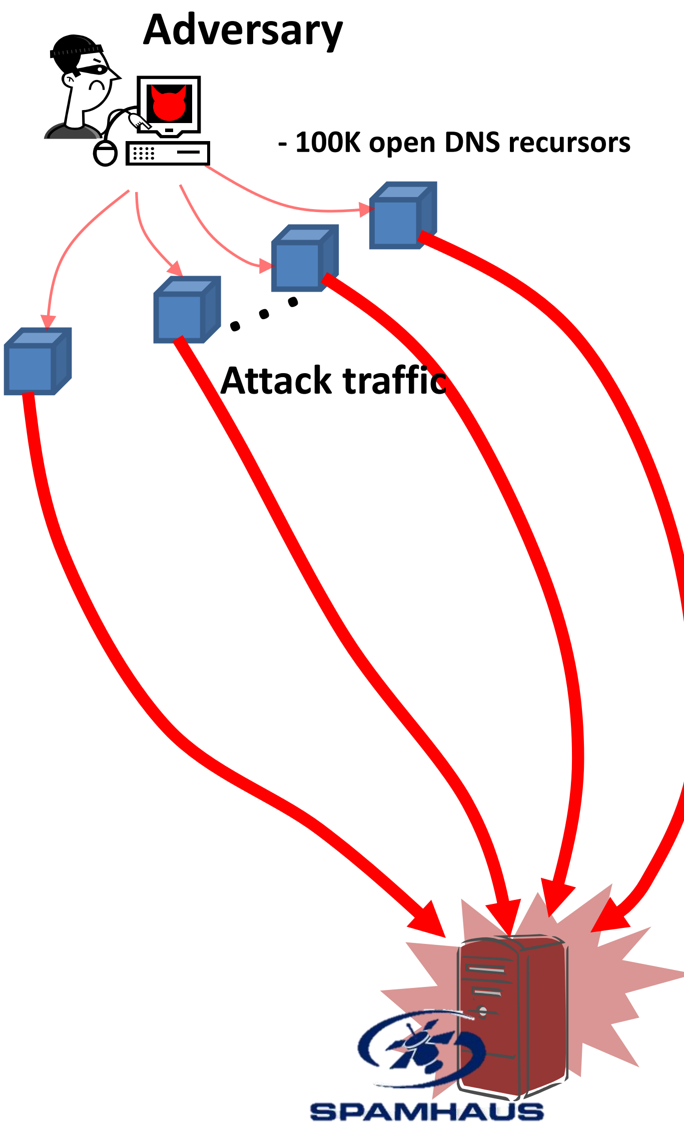

Example: "Spamhaus" Attack (2013)

- Adversary: DDoS -> **1** Spamhaus ***Server***
  3/16 – 3/18: ~ 10 Gbps
  persistent: ~ **2.5 days**
- Spamhaus -> CloudFlare (3/19 – 3/22)
  – **non-scalable: ->** 90-120 Gbps *traffic*
  *is diffused over **N > 20 servers** in* 4 hours

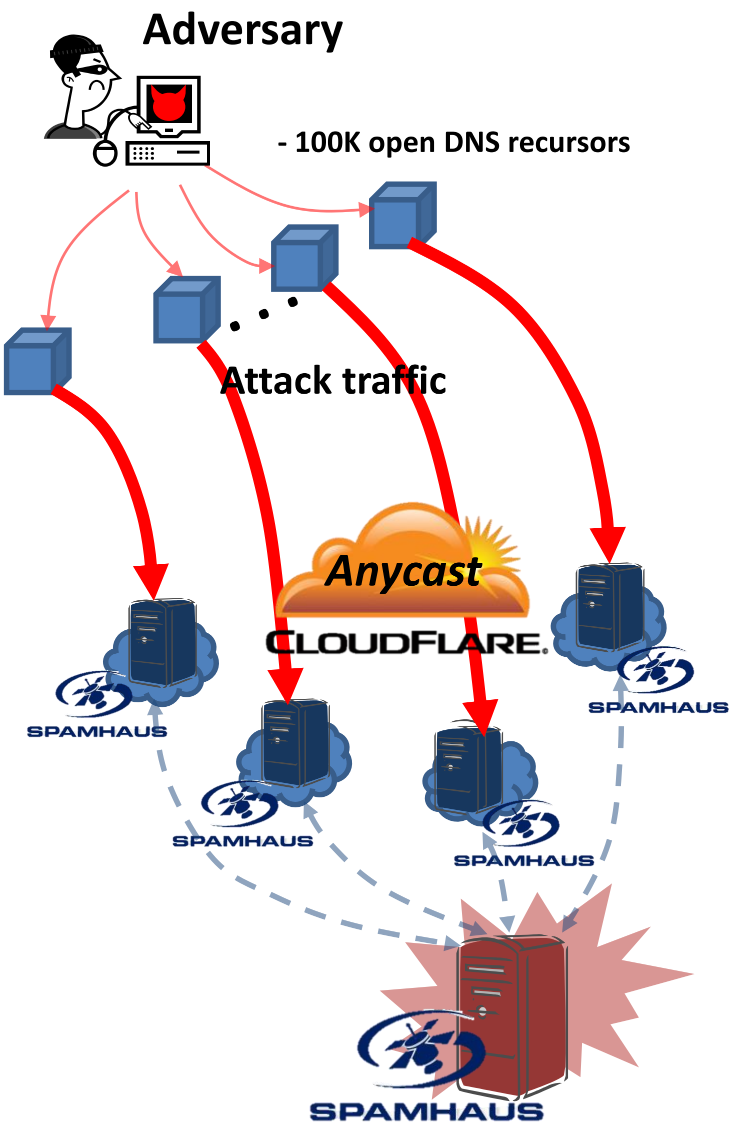

Example: “Spamhaus” Attack (2013)

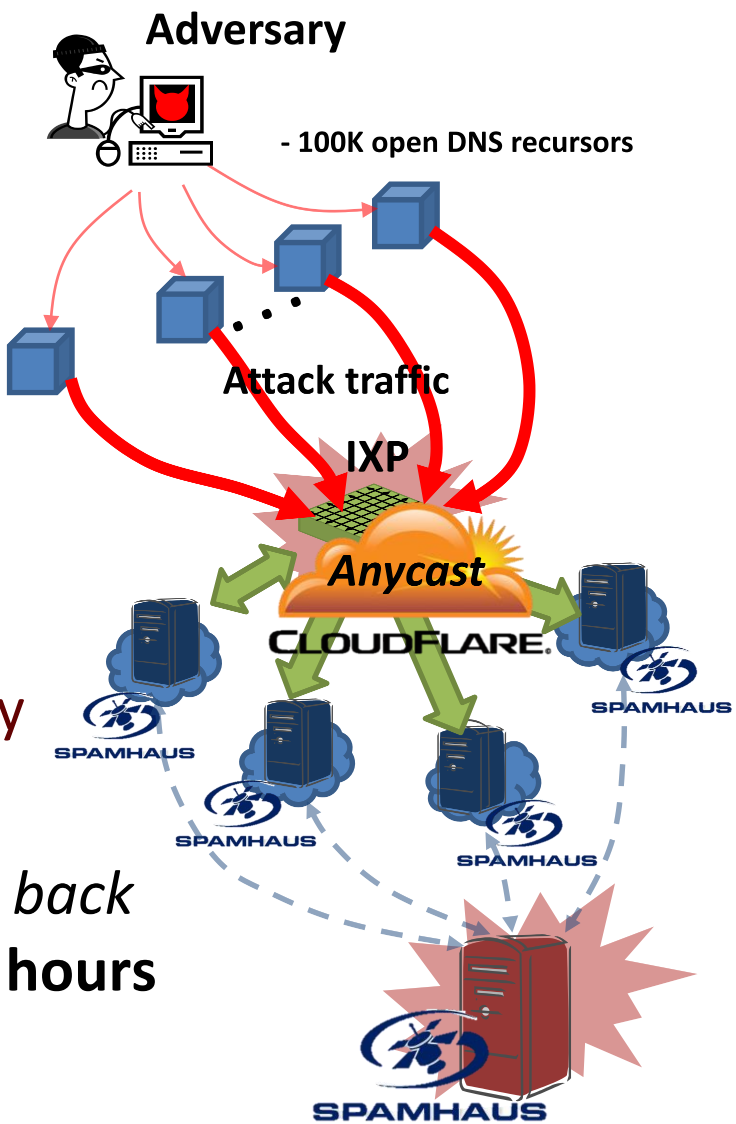

- Adversary: DDoS -> **4 IXPs** (3/23)
  - **scalable**: regionally degraded connectivity
    some disconnection
  - **non-persistent**: *attack detected, pushed back & legitimate traffic re-routed in* ~ **1 - 1.5 hours**

New: The *Crossfire* Attack

*A **link-flooding attack** that degrades/cuts off network connections of **scalable N-server** area **persistently***

➢ Scalable N-Server areas

- **N** = small (e.g., 1 -1000 servers), medium (e.g., all servers in a US state), large (e.g., the West Coast of the US)

➢ Persistent:

- attack traffic is **indistinguishable from legitimate**
  - low-rate, changing sets of flows
- attack is "**moving target" for same N-server area**
  - changes target links before triggering alarms

Definitions

• Target area 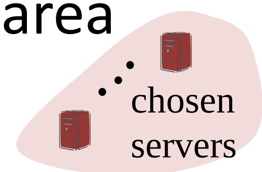	Area containing chosen target servers e.g., an organization, a city, a state, or a country
• Target link 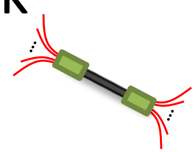	Network link selected for flooding
• Decoy server 	Publicly accessible servers surrounding the target area

1-Link Crossfire

Attack Flows => Indistinguishable from Legitimate

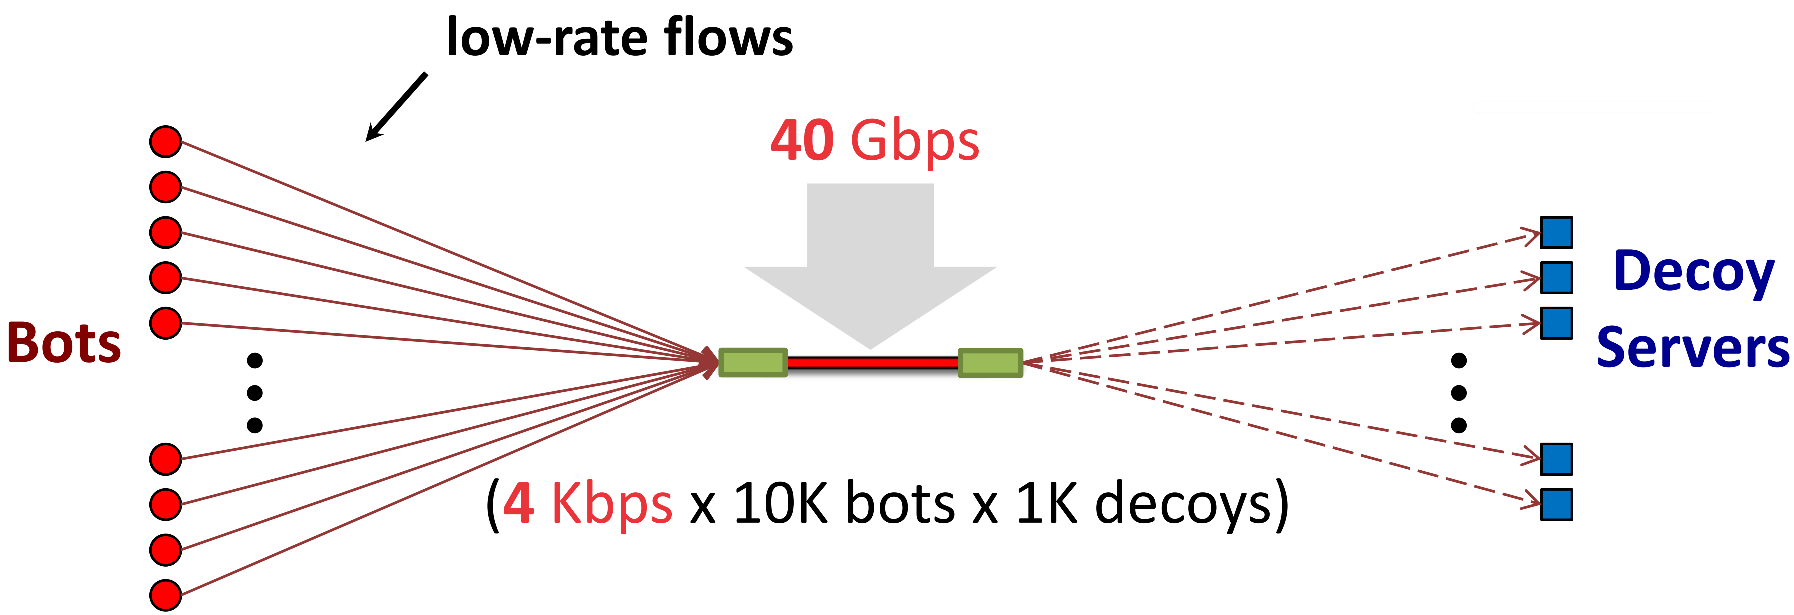

1-Link Crossfire

Attack Flows => Indistinguishable from Legitimate

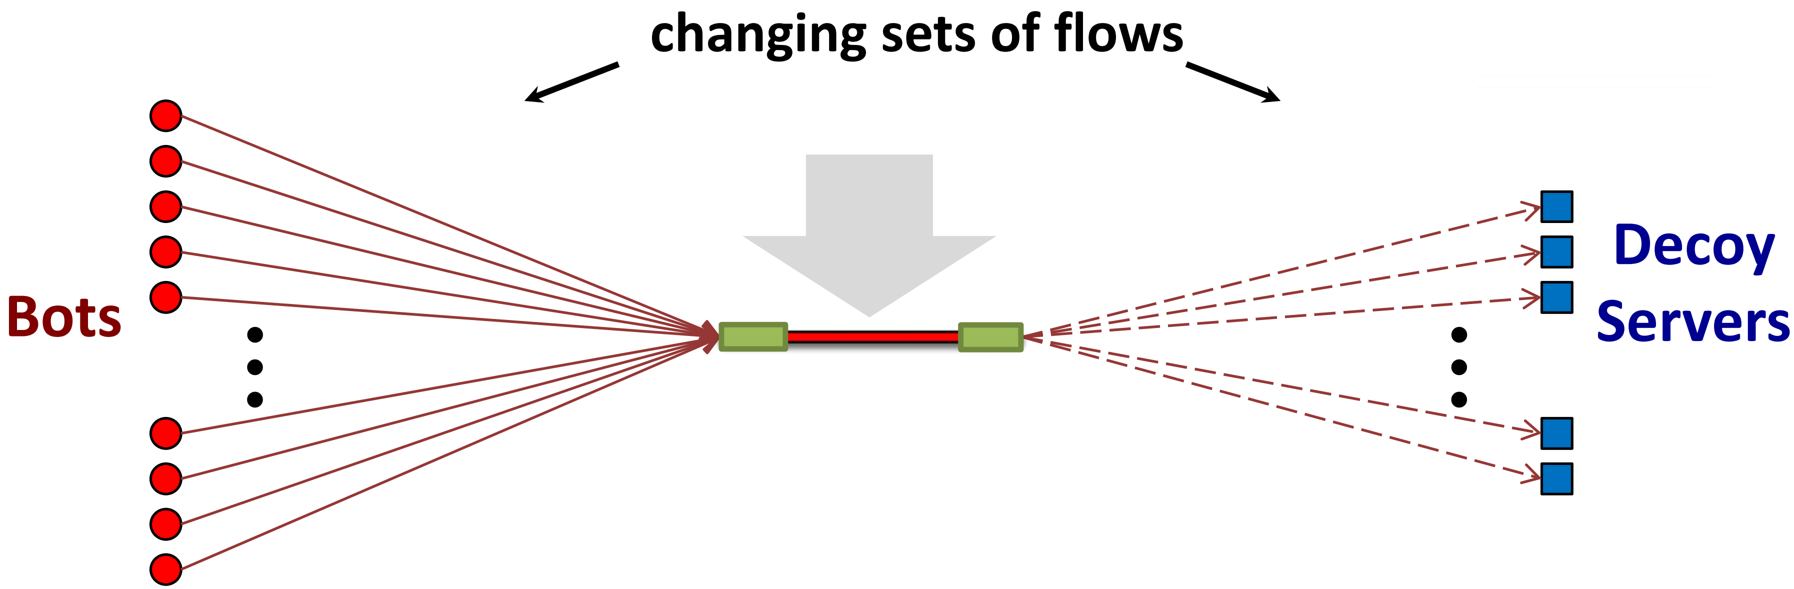

1-Link Crossfire

Attack Flows => Alarms Not Triggered

suspend flows *in* ***t*** **<** $\mathbf{T_{det}}$ **sec** & **resume** later

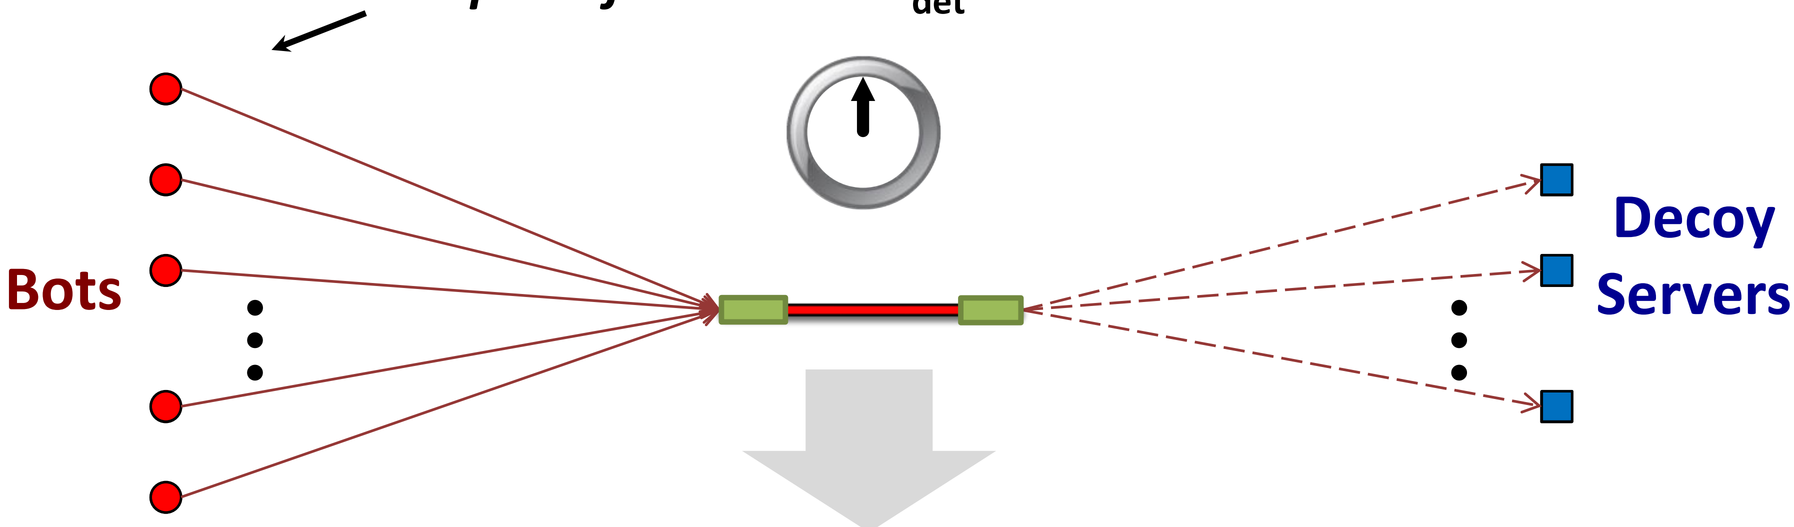

link-failure detection latency, T_{det}

IGP routers: **217 sec**/80 Gbps – **608 sec**/60 Gbps

BGP routers: **1,076 sec**/80Gbps – **11,119 sec**/60 Gbps

t **= 40 – 180 sec => Alarms are Not Triggered**

n-Link Crossfire

- ***n*** **links traversed by a large number of persistent paths to a target area.**

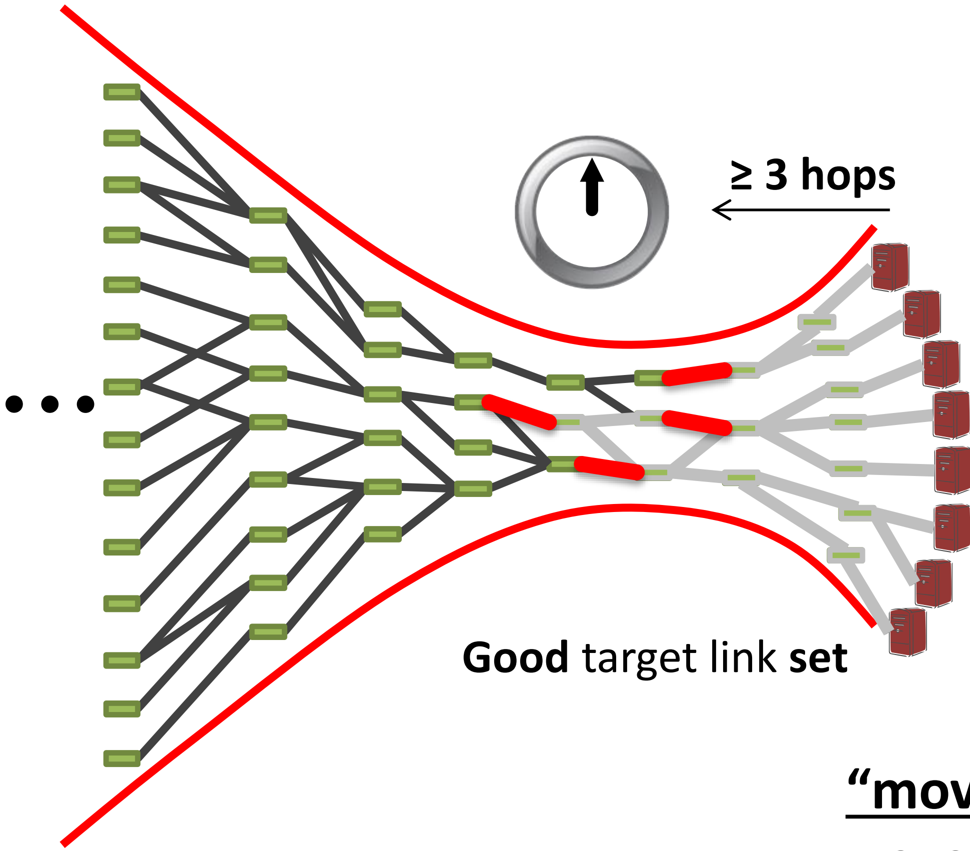

small *n*; e.g., 5 - 15
"Narrow Path Waist"
(observed power law
for Internet route paths)

N servers

"moving targets," same N servers
= suspend-resume *flooding of*
***different* link sets**

n-Link Crossfire

- ***n*** **links traversed by a large number of persistent paths to a target area.**

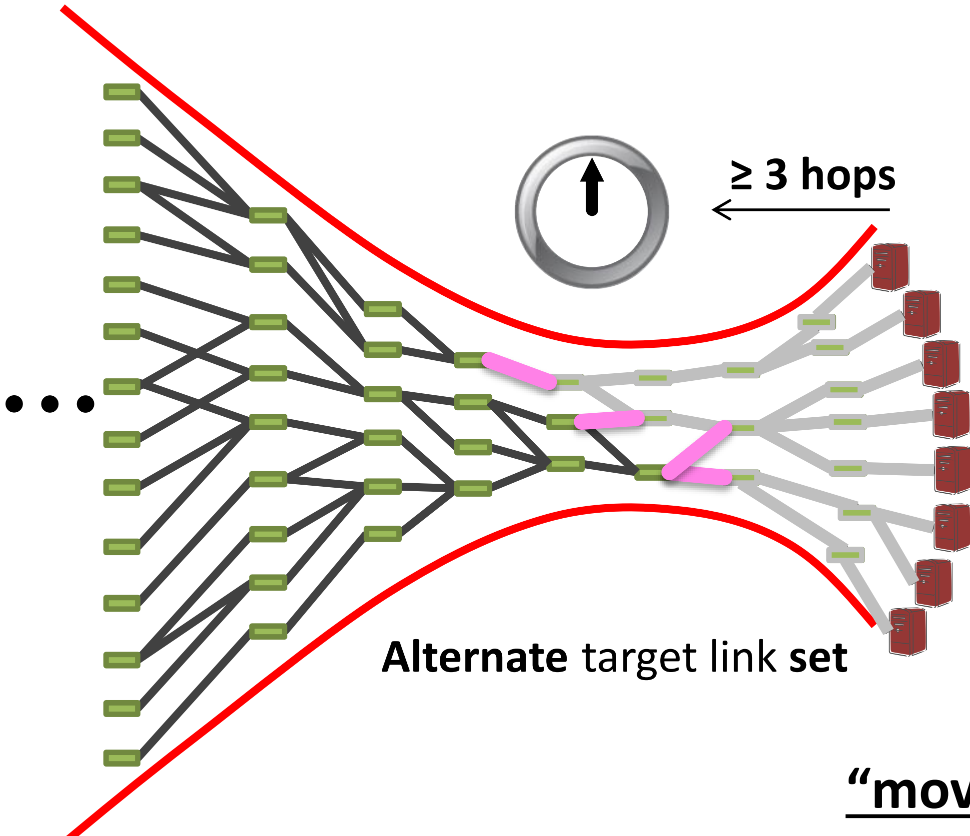

small *n*; e.g., 5 - 15
"Narrow Path Waist"
(observed power law
for Internet route paths)

N servers

"moving targets," same N servers
= suspend-resume *flooding of different* link sets

n-Link Crossfire

- *n* links traversed by a large number of persistent paths to a target area.

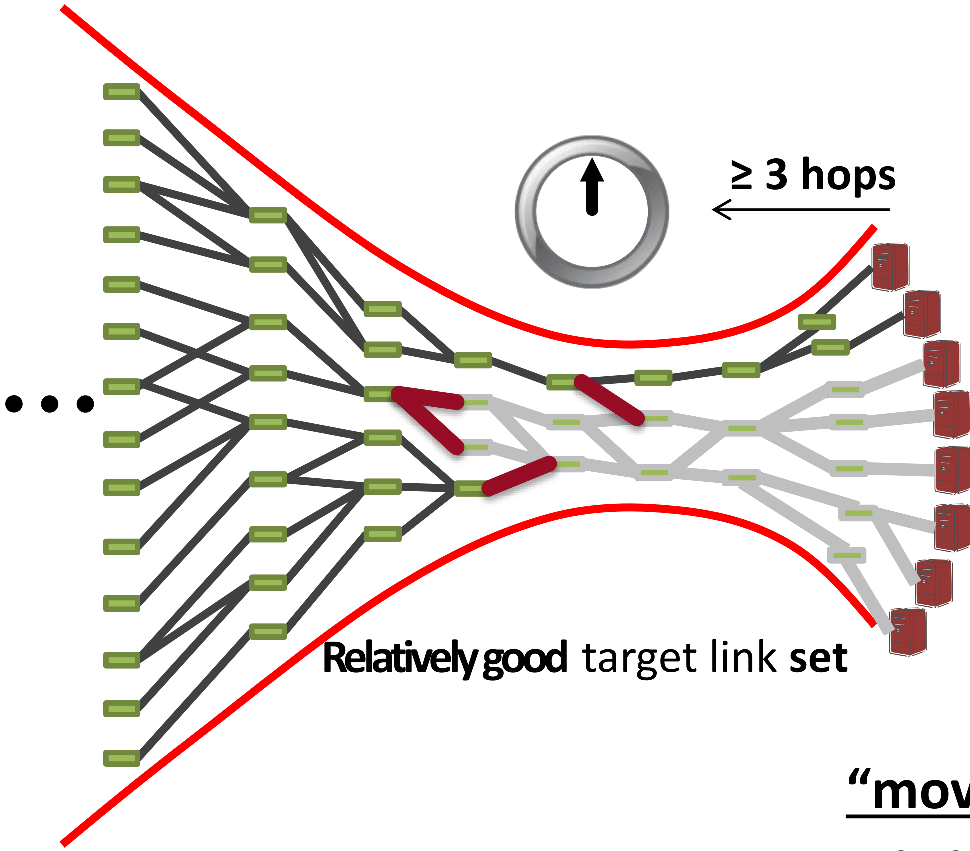

small *n*; e.g., 5 - 15
"Narrow Path Waist"
(observed power law for Internet route paths)

N servers

"moving targets," same N servers = suspend-resume *flooding of different* link sets

Degraded Connectivity

*** Degradation Ratio (target link set) =** $\dfrac{\textbf{\# } \textit{degraded} \text{ bot-to-target area } \textbf{paths}}{\textbf{\# } \textit{all} \text{ bot-to-target area } \textbf{paths}}$

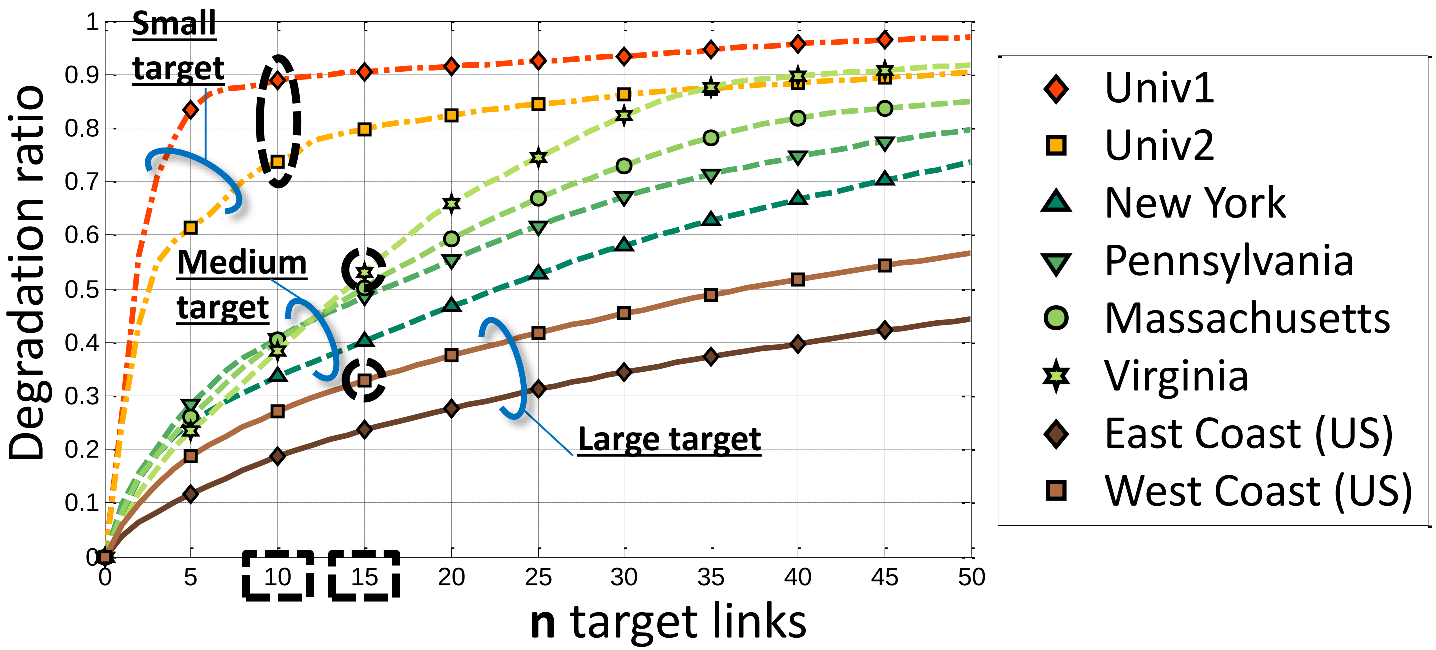

- Flooding *a few* target links causes *high* degradation (DR*)
  - 10 links => DR: 74 – 90% for Univ1 and Univ2
  - 15 links => DR: 53% (33%) for Virginia (West Coast)

Attack Steps

&

Experiments

Attack Step 1: Link-Map Construction

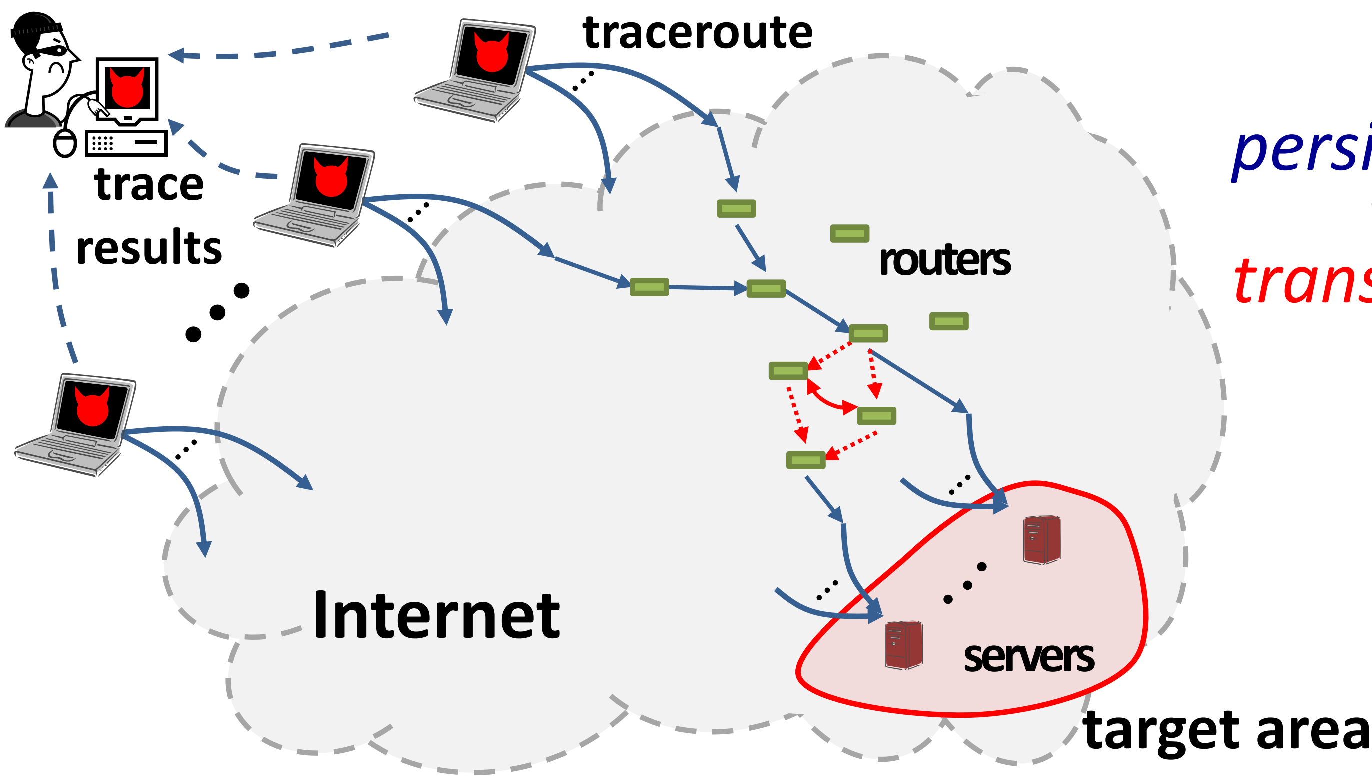

persistent
vs.
transient links

Only *persistent links* are targeted

Attack Step 2: Target-Link Selection

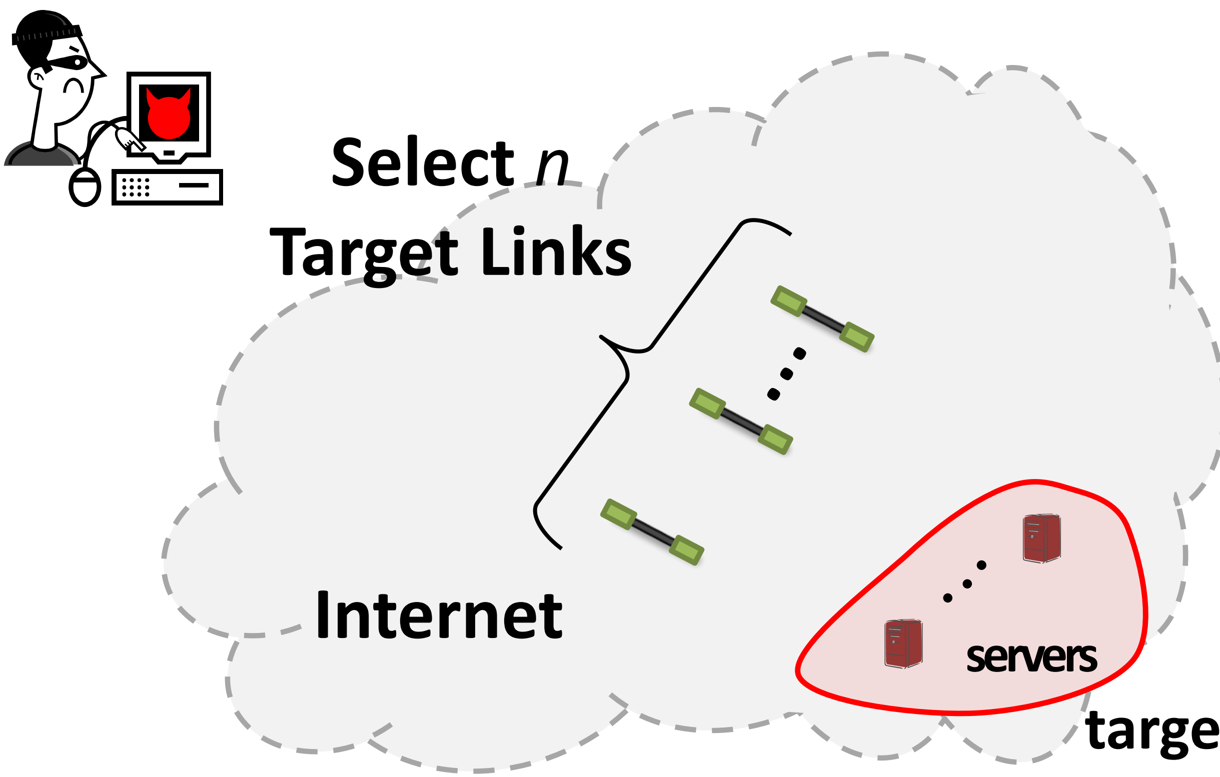

Goal:

Find *n* links whose failure maximizes *DR*

=> maximum coverage problem

Attack Step 3: Bot Coordination

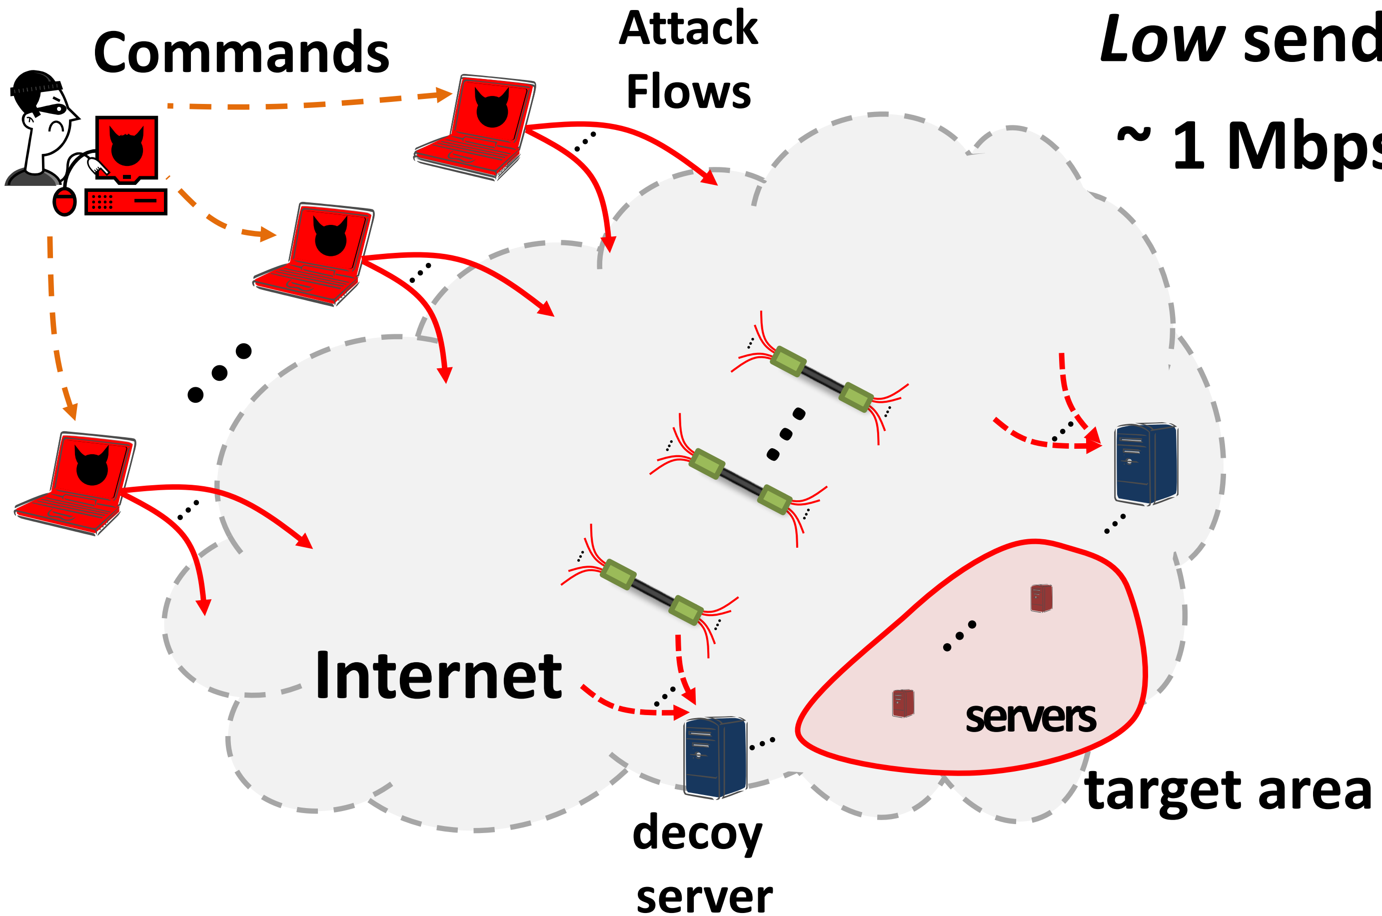

Low send/receive rates

~ 1 Mbps

Experiments

Geographical Distribution of Traceroute Nodes

- 1,072 traceroute nodes
  - 620 PlanetLab nodes + 452 Looking Glass servers

Experiments

Target Areas

Target Areas

small
- Univ1
- Univ2

medium
- New York
- Pennsylvania
- Massachusetts
- Virginia

large
- East Coast
- West Coast

Degraded Connectivity

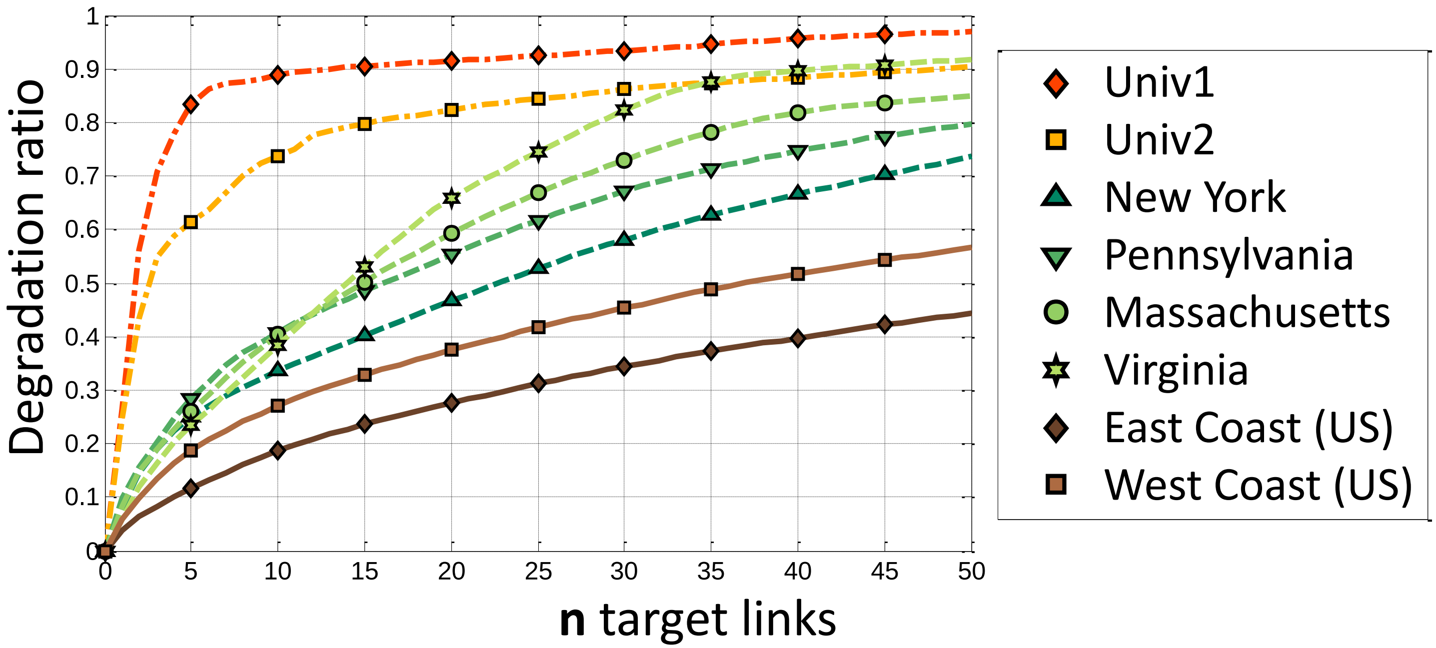

- Flooding *a few* target links causes *high* degradation (DR*)
  - 10 links => DR: 74 – 90% for Univ1 and Univ2
  - 15 links => DR: 53% (33%) for Virginia (West Coast)

Effective Independence of Bot Distribution

Setting:

Experiments using

6 different bot distributions

< Bot distribution on the map >

Result:

No significant difference in attack performance

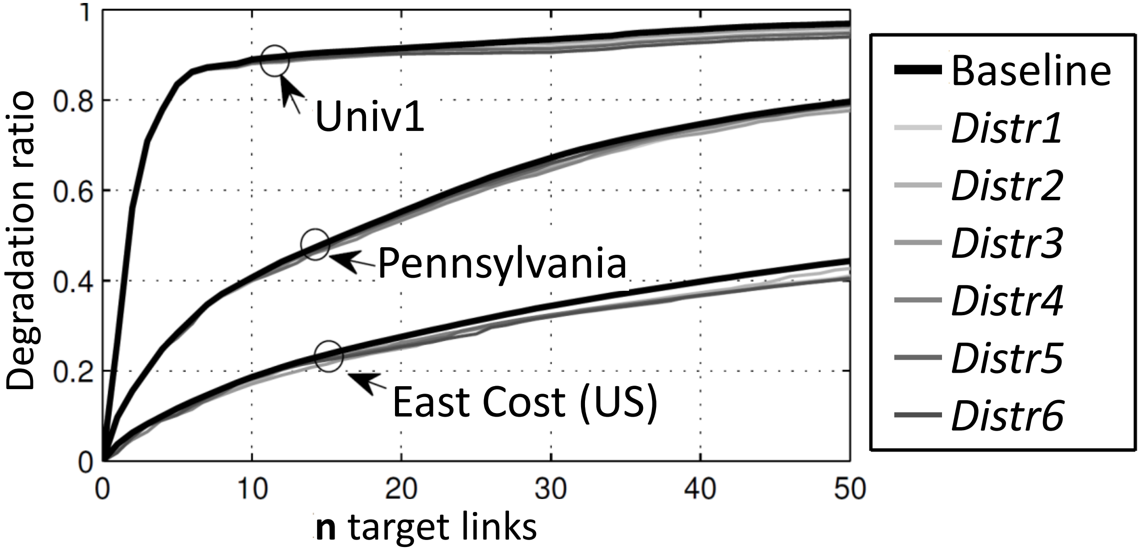

More bots => Lower "Send" Flow Rate

Average rate

when flooding **10 Target Links** against **Pennsylvania**

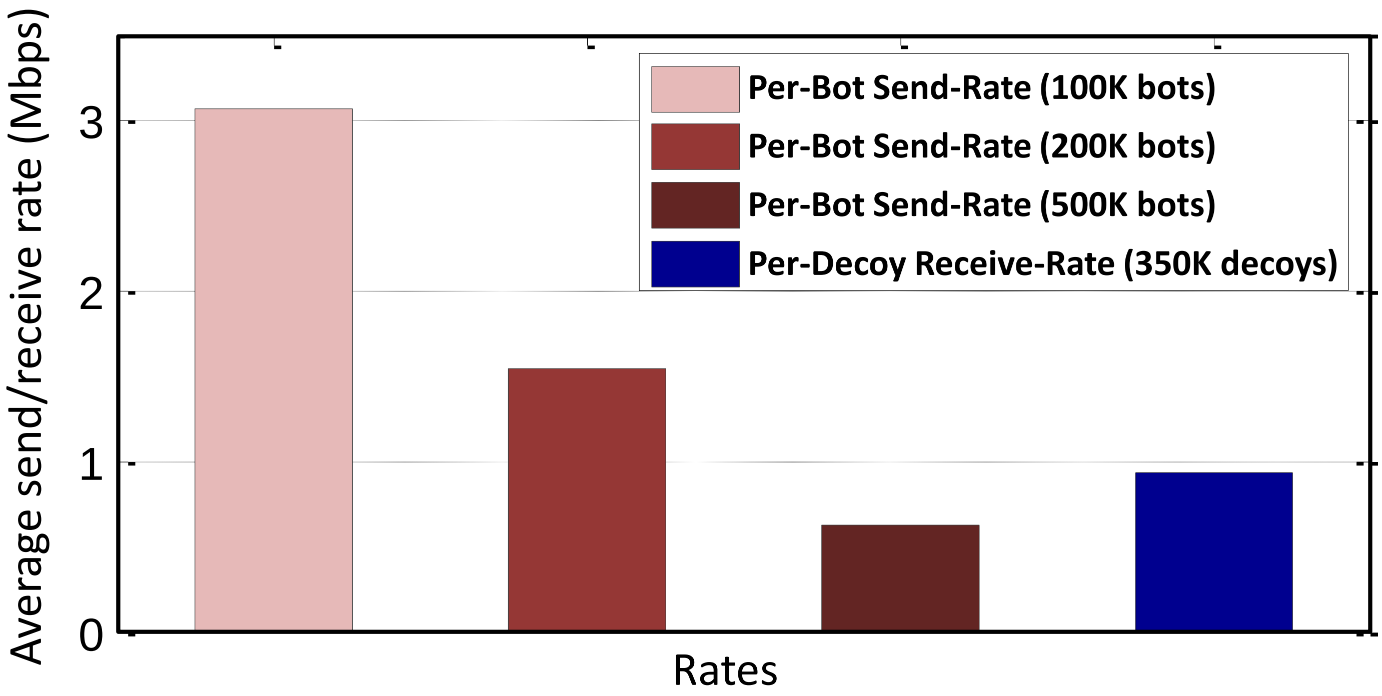

Cost

- Attack bots available from Pay-Per Install (PPI) markets [2011]

Region	Price per thousand bots
US / UK	$100 - $180
Continental Europe	$20 - $60
Rest of the world	< $10

  – 10 target link flooding
    » 500 K bots => $46K
    » 100 K bots => $9K

- State-/corporate-sponsored attacks use 10 – 100 x more bots
- Zero cost; e.g., harvest 100 – 500 K bots for 10 links

Crossfire vs. Other Attacks

Design Goal	Old DDoS	Coremelt (2009)	"Spamhaus" Attack (2013)	Crossfire (2013)
Scalable choice of **N server targets**	✗	Not a Goal	✗	✓
Bot distribution **independence**	✓	✗	Not a Goal	✓
Indistinguishability from **Legitimate** flows	✗	✓	✗	✓
Reliance on **wanted flows** only	✗	✓	✗	✗
Persistence	HIGH	Low	Low	HIGH

Possible Countermeasures

- Any countermeasure must address (at least one of)
  i. the *existence* of the *"narrow path waist"*
  ii. slow network & ISP *reaction*
- *Cooperation among multiple ISPs* becomes necessary for detection
- Application-layer *overlays* can route around flooded links

- Additional measures
  - Preemptive or retaliatory *disruption of bot markets*
  - International agreements regarding prosecution of telecommunication-infrastructure attacks

Conclusion

- New DDoS attack: the Crossfire attack
  - Scalable **&** Persistent
- Internet-scale experiments
  - Feasibility of the attack
  - High impact with low cost
- Generic Countermeasures
  - Characterization of possible solutions

Questions?

Min Suk Kang

minsukkang@cmu.edu