CS8803 - EMS

Class 02 — Denial of Service: Backscatter

Paul Pearce

Please wear a mask in this classroom

Welcome!

Overview of Today

- Course logistics and introductions
- Summary of the "Inferring Internet Denial of Service Activity
  - Moore, Voelker and Savage, USENIX Security 2001
- What do you think?
- Guided Discussion
- Action Items

Logistics

Introductions

- Me
  - Paul Pearce
  - Assistant Professor, School of Cybersecurity and Privacy, School of Computer Science
  - PhD UC Berkeley Computer Science, 2018
    - Advised by Vern Paxson
    - Worked closely with UC San Diego and Princeton
  - Spent a year as a visiting researcher Facebook
  - MS and BS also UC Berkeley
    - (Go Bears!)

Introductions

- You
  - Your name
  - Your program
  - How far into your program
  - Have you taken a graduate security course before?
  - An interesting fact (if you choose to share)
  - What do you hope to get out of this class?

Inferring Internet Denial-of-Service Activity

Network Behaviors

Packet sent	Response from victim
TCP SYN (to open port)	TCP SYN/ACK
TCP SYN (to closed port)	TCP RST (ACK)
TCP ACK	TCP RST (ACK)
TCP DATA	TCP RST (ACK)
TCP RST	no response
TCP NULL	TCP RST (ACK)
ICMP ECHO Request	ICMP Echo Reply
ICMP TS Request	ICMP TS Reply
UDP pkt (to open port)	protocol dependent
UDP pkt (to closed port)	ICMP Port Unreach
...	...

Table 1: A sample of victim responses to typical attacks.

Spoofing and Backscatter

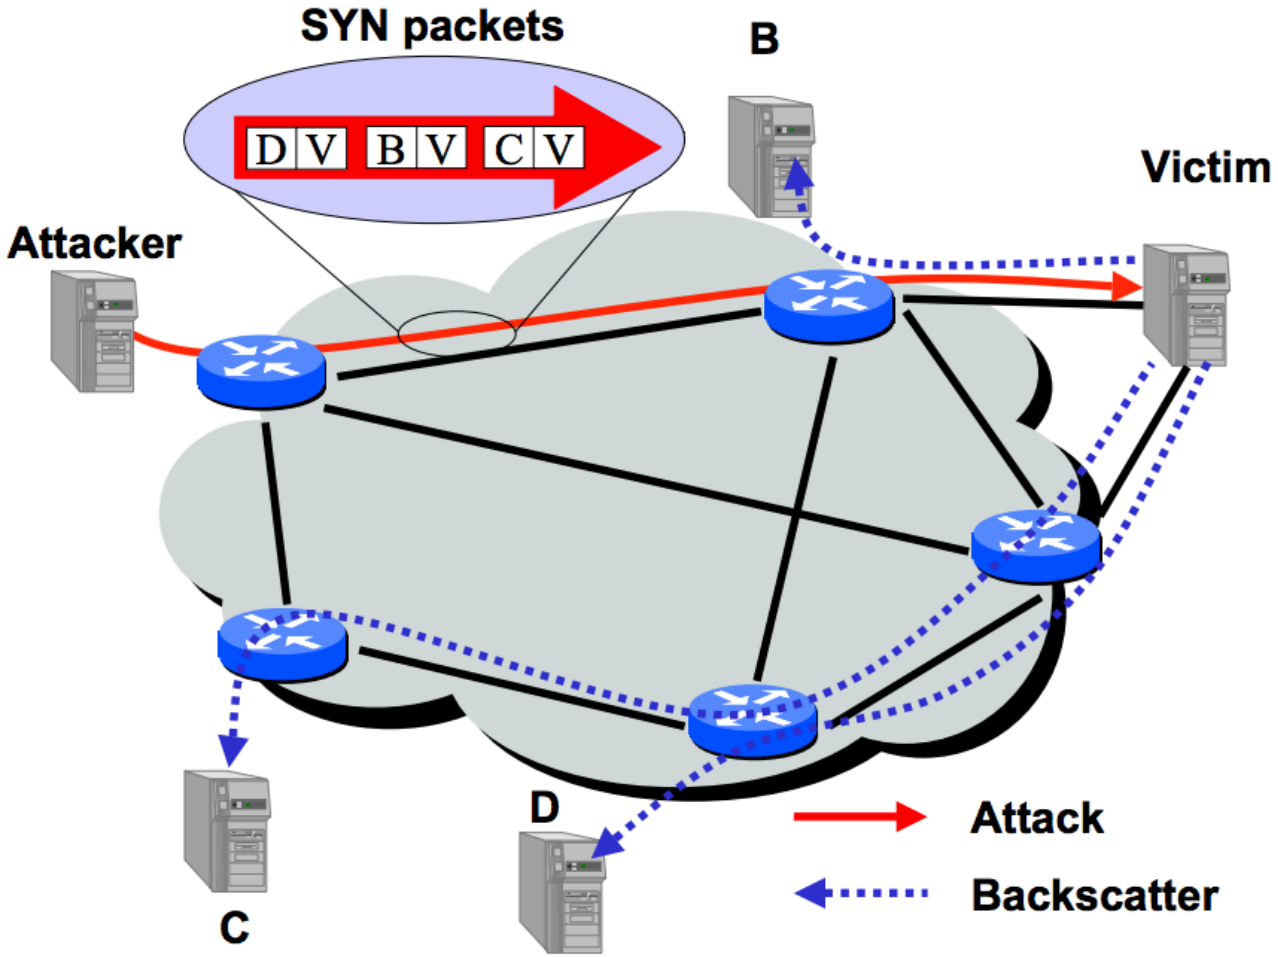

- Expectation of observing an attack
  - E(X) = N * M * I / (2^32)

- Attack rate R
  - R >= R' * 2^32 / N

Results

	Trace-1	Trace-2	Trace-3
Dates (2001)	Feb 01 – 08	Feb 11 – 18	Feb 18 – 25
Duration	7.5 days	6.2 days	7.1 days
Flow-based Attacks:			
Unique victim IPs	1,942	1,821	2,385
Unique victim DNS domains	750	693	876
Unique victim DNS TLDs	60	62	71
Unique victim network prefixes	1,132	1,085	1,281
Unique victim Autonomous Systems	585	575	677
Attacks	4,173	3,878	4,754
Total attack packets	50,827,217	78,234,768	62,233,762
Event-based Attacks:			
Unique victim IPs	3,147	3,034	3,849
Unique victim DNS domains	987	925	1,128
Unique victim DNS TLDs	73	71	81
Unique victim network prefixes	1,577	1,511	1,744
Unique victim Autonomous Systems	752	755	874
Attack Events	112,457	102,204	110,025
Total attack packets	51,119,549	78,655,631	62,394,290

Table 2: Summary of backscatter database.

Time series data

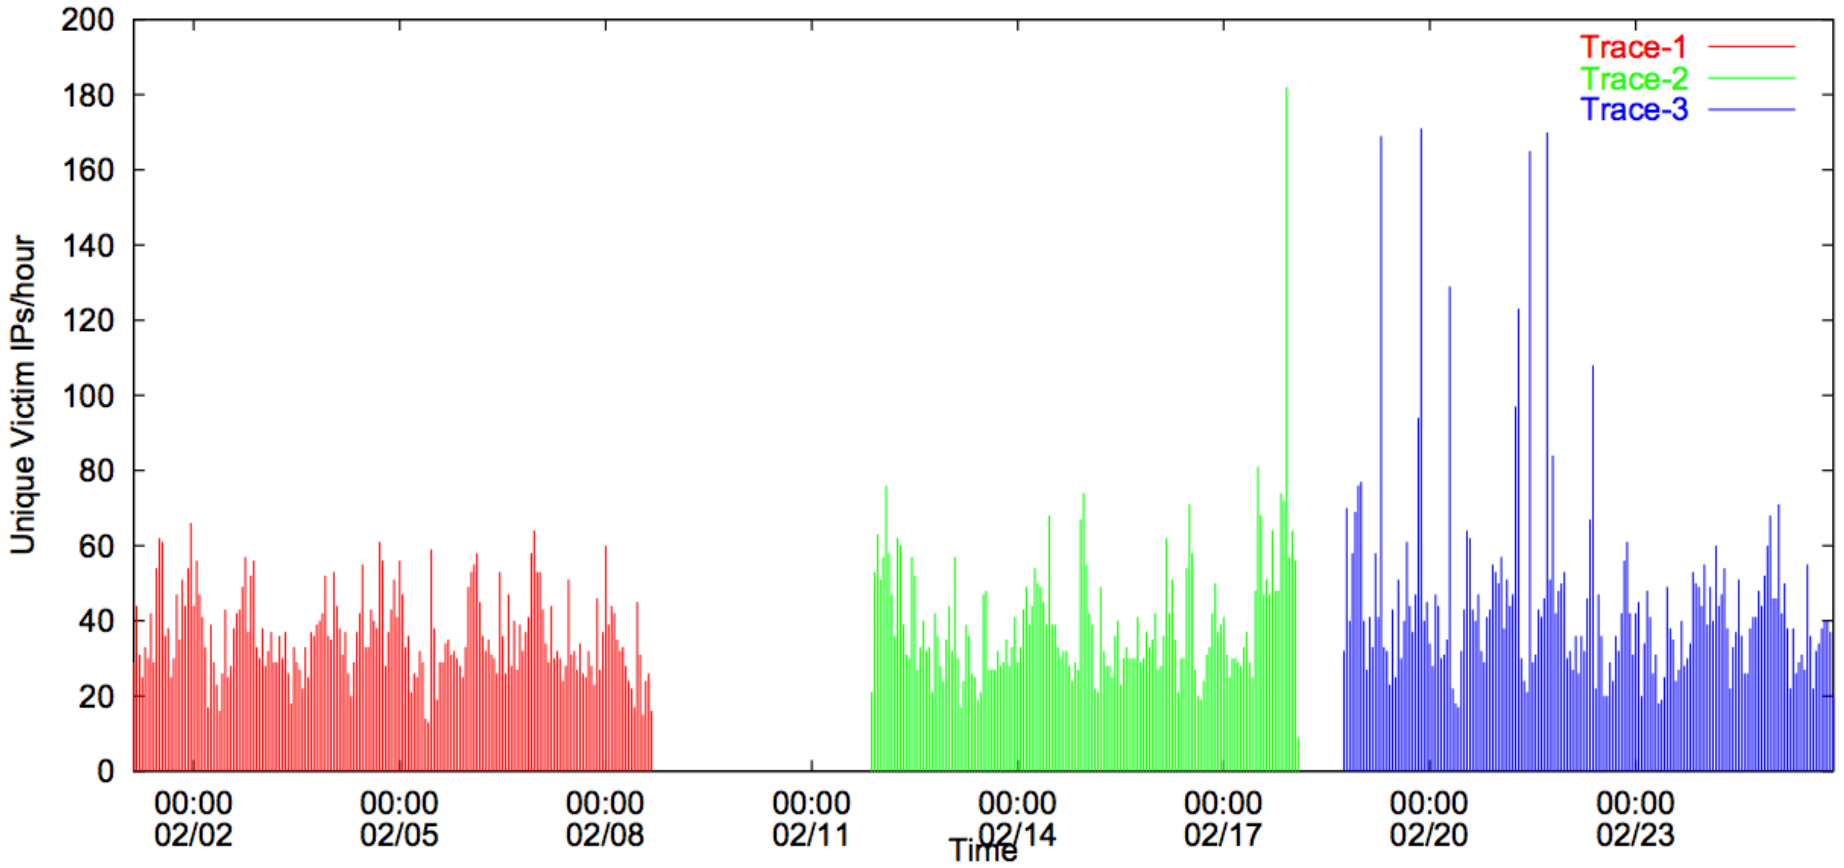

Figure 3: Estimated number of attacks per hour as a function of time (UTC).

Packet types

Kind	Trace-1		Trace-2		Trace-3	
	Attacks	Packets (k)	Attacks	Packets (k)	Attacks	Packets (k)
TCP (RST ACK)	2,027 (49)	12,656 (25)	1,837 (47)	15,265 (20)	2,118 (45)	11,244 (18)
ICMP (Host Unreachable)	699 (17)	2,892 (5.7)	560 (14)	27,776 (36)	776 (16)	19,719 (32)
ICMP (TTL Exceeded)	453 (11)	31,468 (62)	495 (13)	32,001 (41)	626 (13)	22,150 (36)
ICMP (Other)	486 (12)	580 (1.1)	441 (11)	640 (0.82)	520 (11)	472 (0.76)
TCP (SYN ACK)	378 (9.1)	919 (1.8)	276 (7.1)	1,580 (2.0)	346 (7.3)	937 (1.5)
TCP (RST)	128 (3.1)	2,309 (4.5)	269 (6.9)	974 (1.2)	367 (7.7)	7,712 (12)
TCP (Other)	2 (0.05)	3 (0.01)	0 (0.00)	0 (0.00)	1 (0.02)	0 (0.00)

Table 3: Breakdown of response protocols.

Uniform Assumption

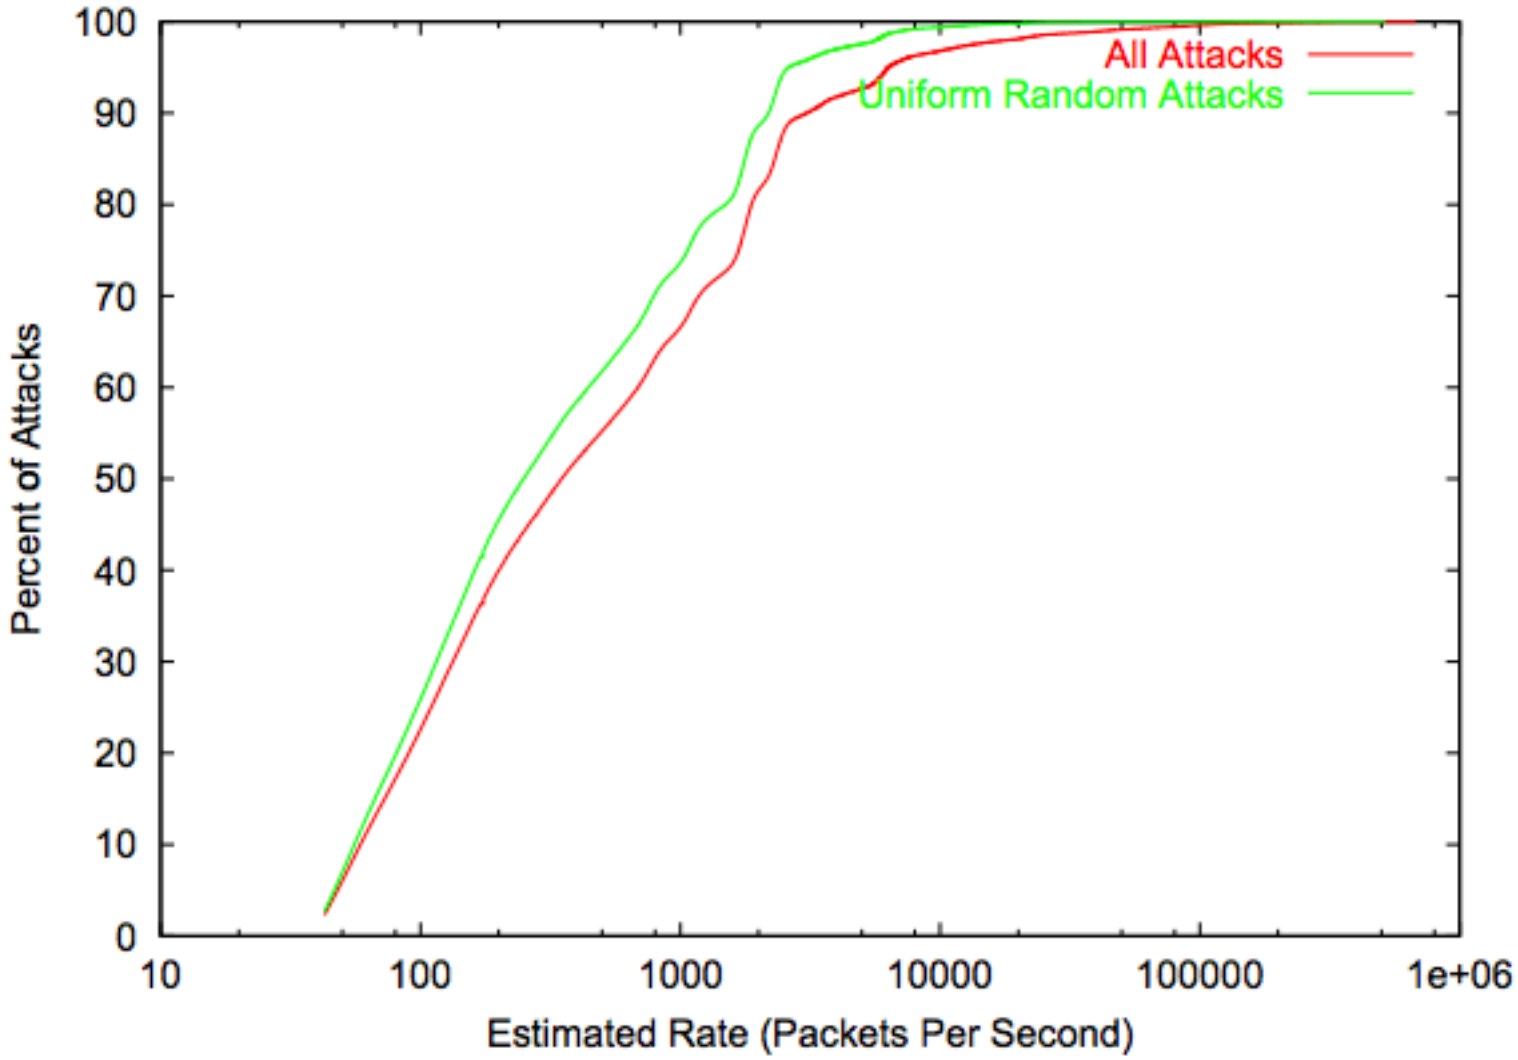

Attack Duration

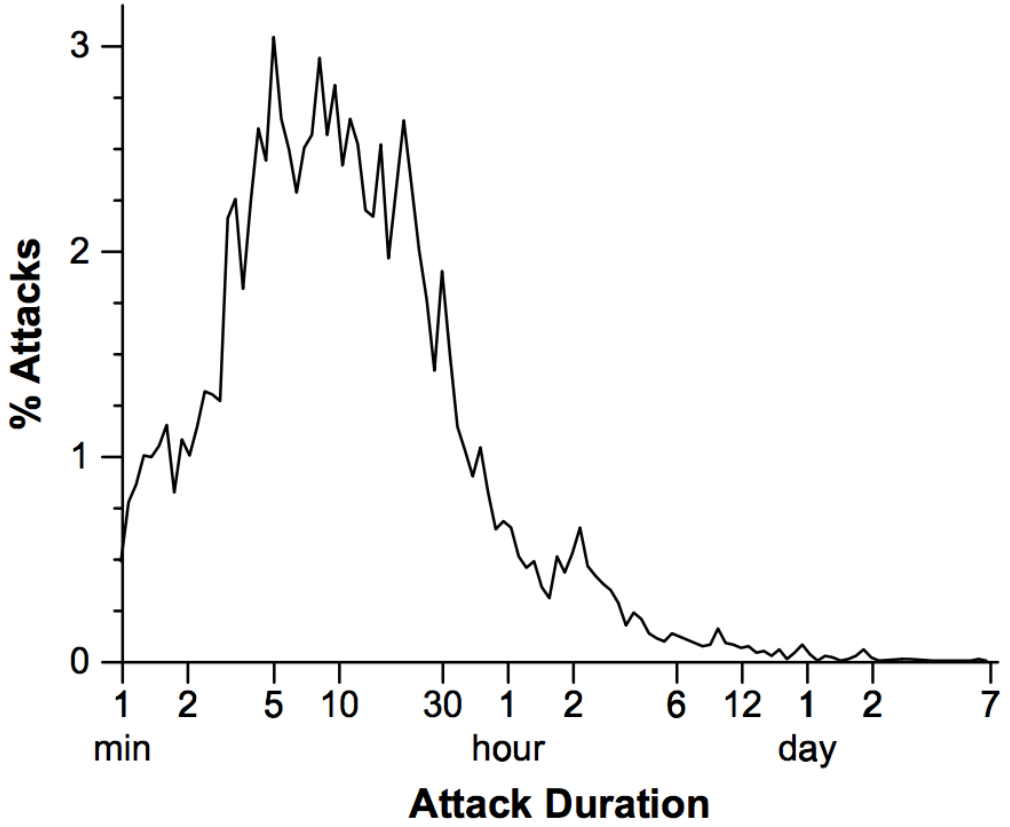

Figure 6: Probability density of attack durations.

Validation

Next, we were able to duplicate a portion of our analysis using data provided by Vern Paxson taken from several University-related networks in Northern California. This new dataset covers the same period, but only detects TCP backscatter with the SYN and ACK flags set. The address space monitored was also much smaller, consisting of three /16 networks ($\frac{3}{65536}$'s of the total IP address space). For 98% of the victim IP addresses recorded in this smaller dataset, we find a corresponding record at the same time in our larger dataset. We can think of no other mechanism other than backscatter that can explain such a close level of correspondence.

Your Thoughts

Guided Discussion

Action Items

- Summary and questions for class 3 (Mirai botnet) will go online tonight
  - Questions for Tuesday class will go online by the end of Thursdays (at the latest)
  - Questions for Thursday will go online by the end of Sunday (at the latest)

- Summary and questions for class 3 are due Monday at noon
  - How's that time sound for you all?

- Discussion lead signup forms will go online at 5pm (~1.75hr from now)
  - A Canvas note will be sent with the URL

Thank You

Questions?

Paul Pearce
https://cc.gatech.edu/~pearce/