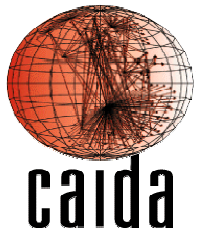

Worldwide Detection of Denial of Service (DoS) Attacks

David Moore,

Geoff Voelker and Stefan Savage

August 15, 2001

dmoore @ caida.org

www.caida.org

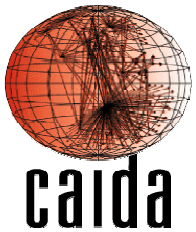

Outline

- The Backscatter Analysis Technique
- Observations and Results
- Validation
- Conclusions

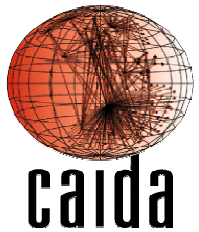

Key Idea

- Backscatter analysis provides *quantitative data* for a **global view** on DoS activity using **local monitoring**

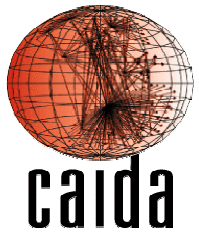

Backscatter Analysis Technique

- Flooding-style DoS attacks
  - e.g. SYN flood, ICMP flood
- Attackers spoof source address **randomly**
  - True of all major attack tools
  - i.e. not SMURF or reflector attack
- Victims, in turn, respond to attack packets
- Unsolicited responses (*backscatter*) equally distributed across IP space
- Received backscatter is **evidence** of an attacker elsewhere

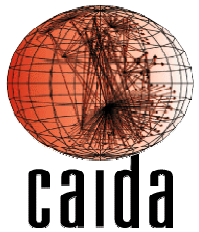

Backscatter Example: Responses Monitored

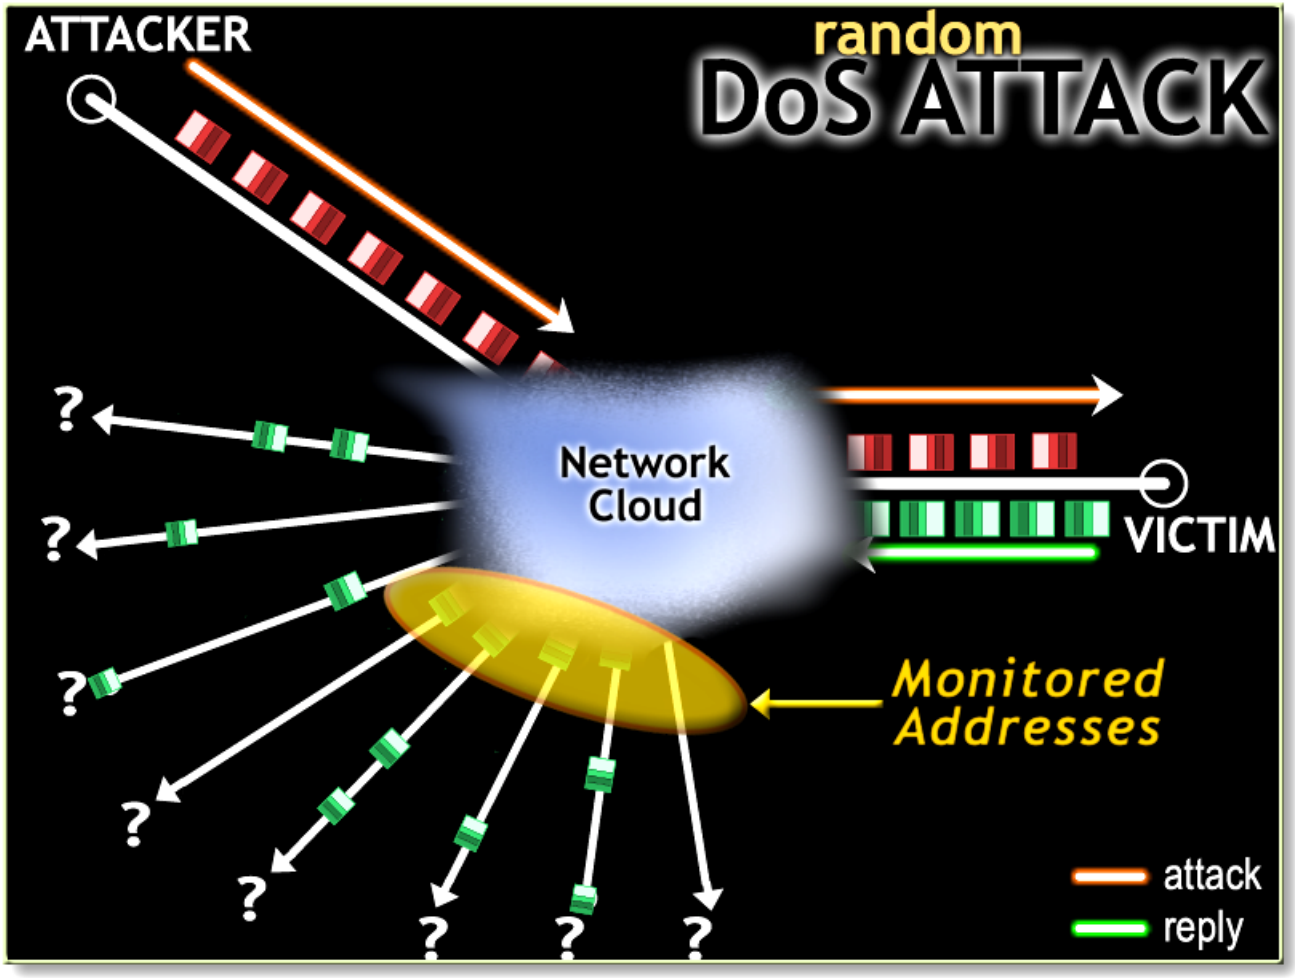

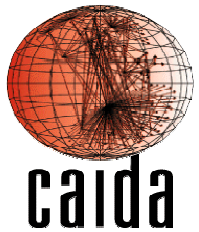

Backscatter analysis

- Monitor block of *n* IP addresses
- Expected # of backscatter packets given an attack of *m* packets:

$$E(X) = \frac{nm}{2^{32}}$$

- Extrapolated attack rate R’ is a function of measured backscatter rate R:

$$R \geq R' \frac{2^{32}}{n}$$

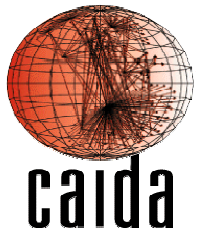

Assumptions and biases

- *Address uniformity*
  - Ingress filtering, reflectors, etc. cause us to **underestimate** # of attacks
  - Can bias rate estimation (can we test uniformity?)
- *Reliable delivery*
  - Packet losses, server overload & rate limiting cause us to **underestimate** attack rates/durations
- *Backscatter hypothesis*
  - Can be biased by purposeful unsolicited packets
    - Port scanning (minor factor at worst in practice)
  - Do we detect backscatter at multiple sites?

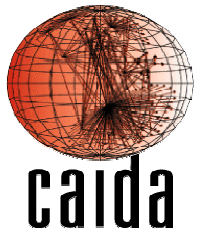

Identifying attacks

- Flow-based analysis (categorical)
  - Keyed on victim IP address and protocol
  - Flow duration defined by explicit parameters (min. threshold, timeout)

- Event-based analysis (intensity)
  - Attack event: backscatter packets from IP address in 1 minute window
  - No notion of attack duration or “kind”

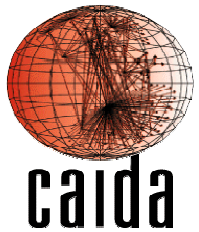

Results

- Attack Breakdown
  - Attacks over Time
  - Protocol Characterization
  - Duration
  - Rate
- Victim Characterization
  - By hostname
  - By TLD

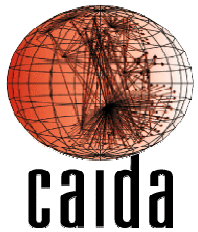

Attack breakdown (three weeks in February)

	Week1	Week2	Week3
Attacks	4173	3878	4754
Victim IPs	1942	1821	2385
Victim prefixes	1132	1085	1281
Victim ASes	585	575	677
Victim DNS domains	750	693	876
Victim DNS TLDs	60	62	71

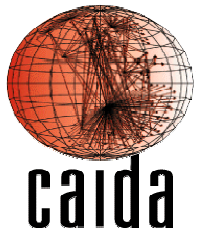

Attacks over time

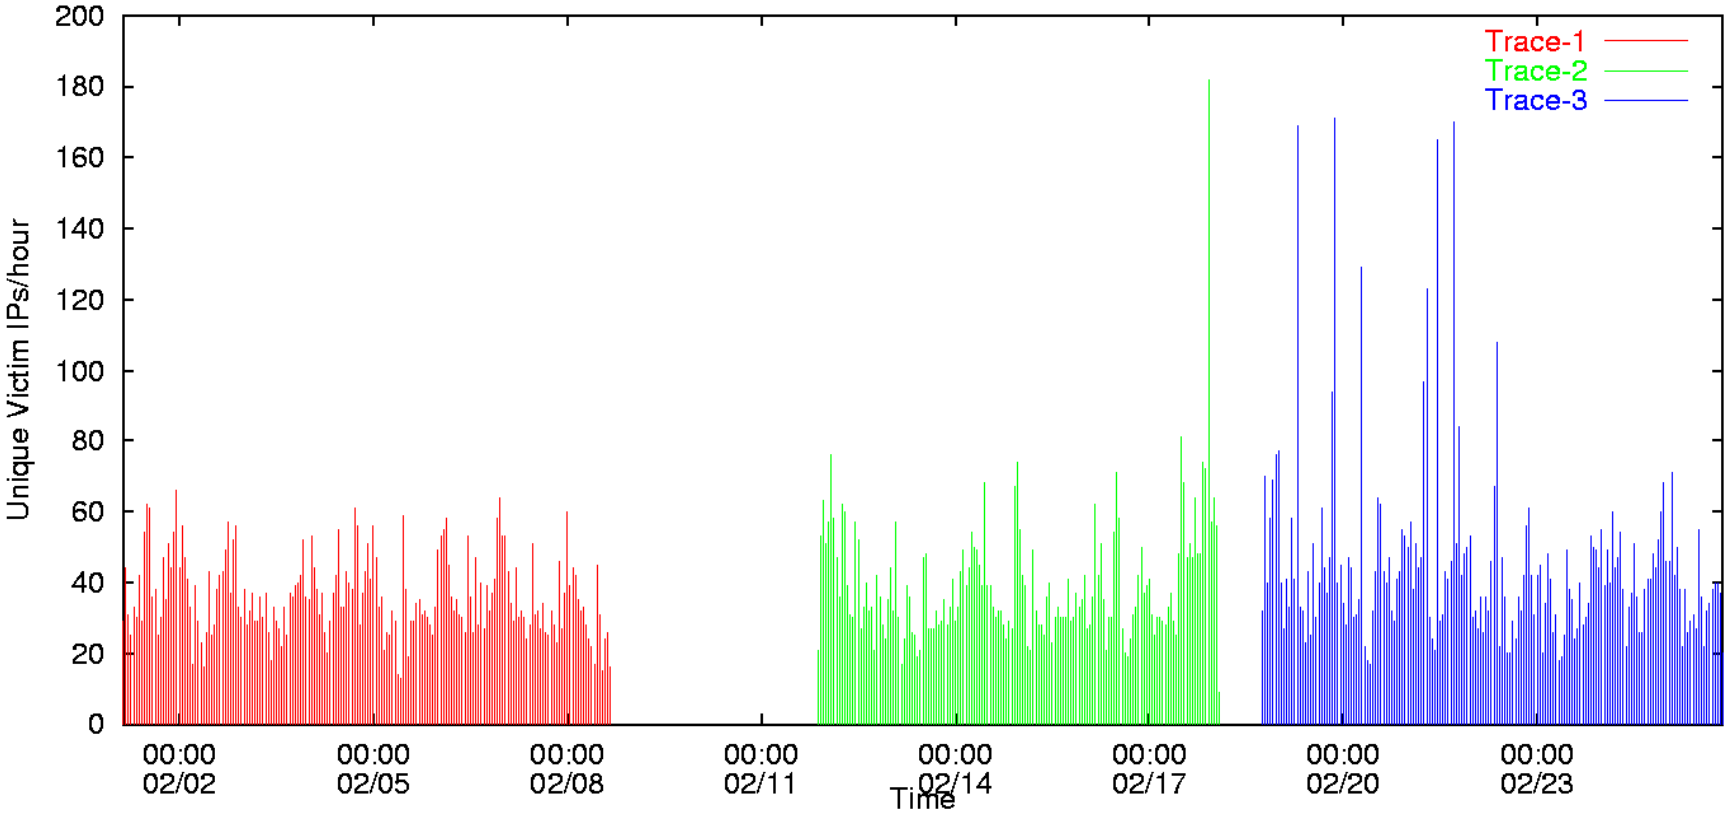

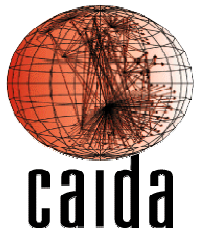

Attack characterization

- Protocols
  - Mostly TCP (90-94% attacks), but a few large ICMP floods (up to 43% of packets)
  - Some evidence of ISP “blackholing” (ICMP host unreachable)
- Services
  - Most attacks on multiple ports (~80%)
  - A few services (HTTP, IRC) singled out

Attack duration distribution

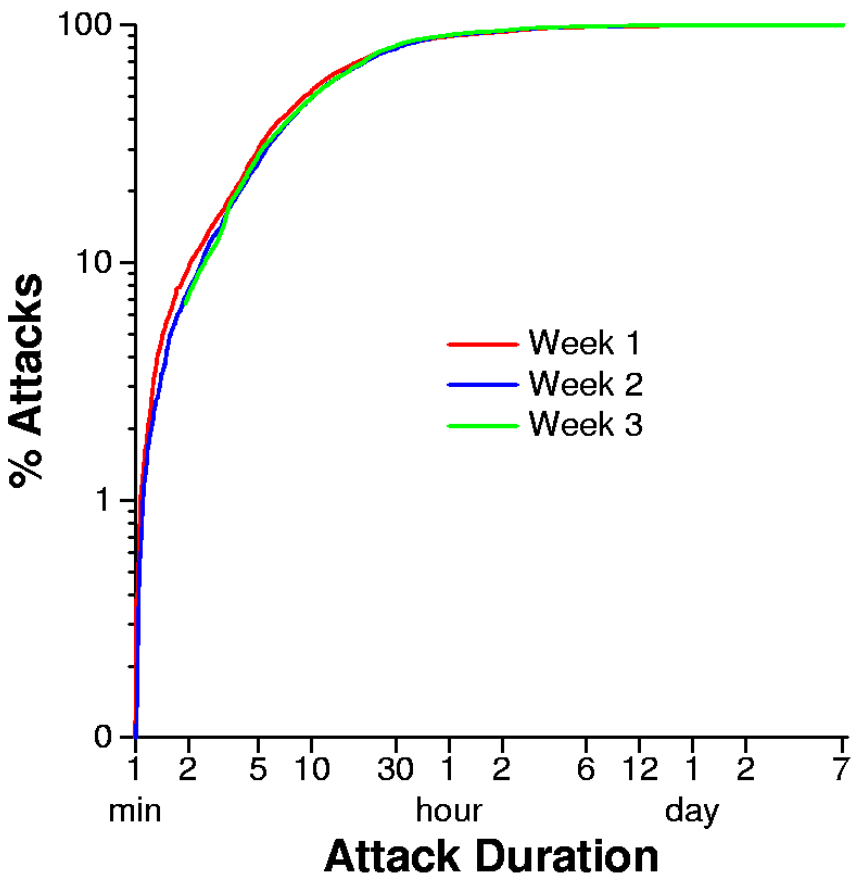

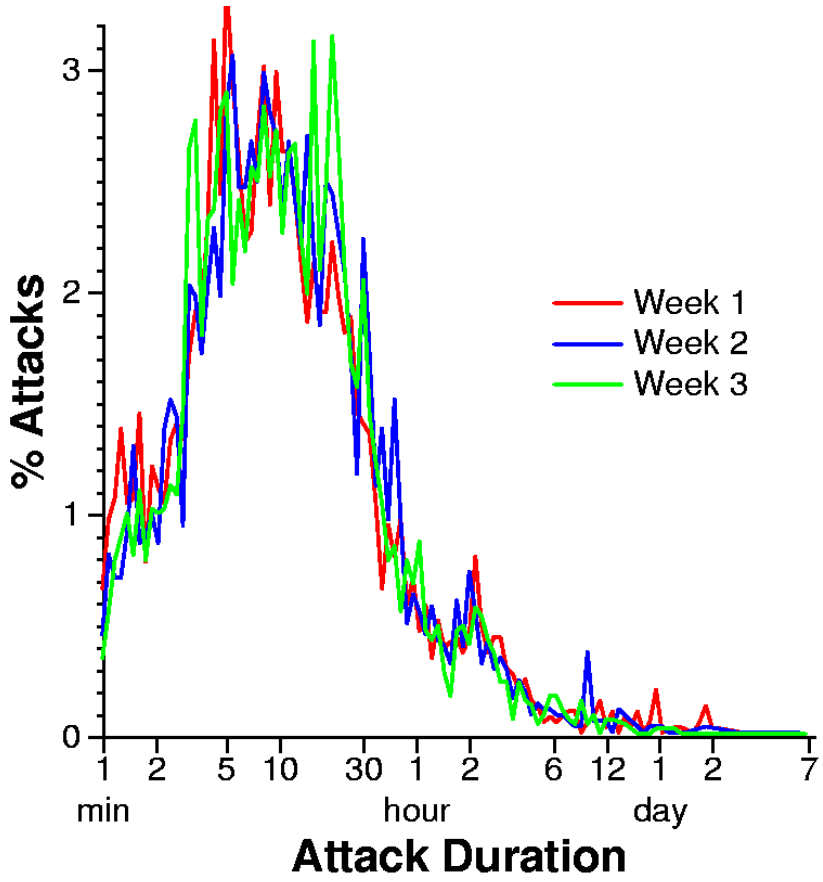

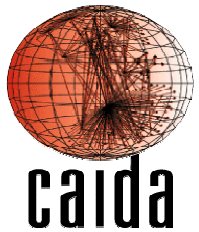

Attack rate distribution

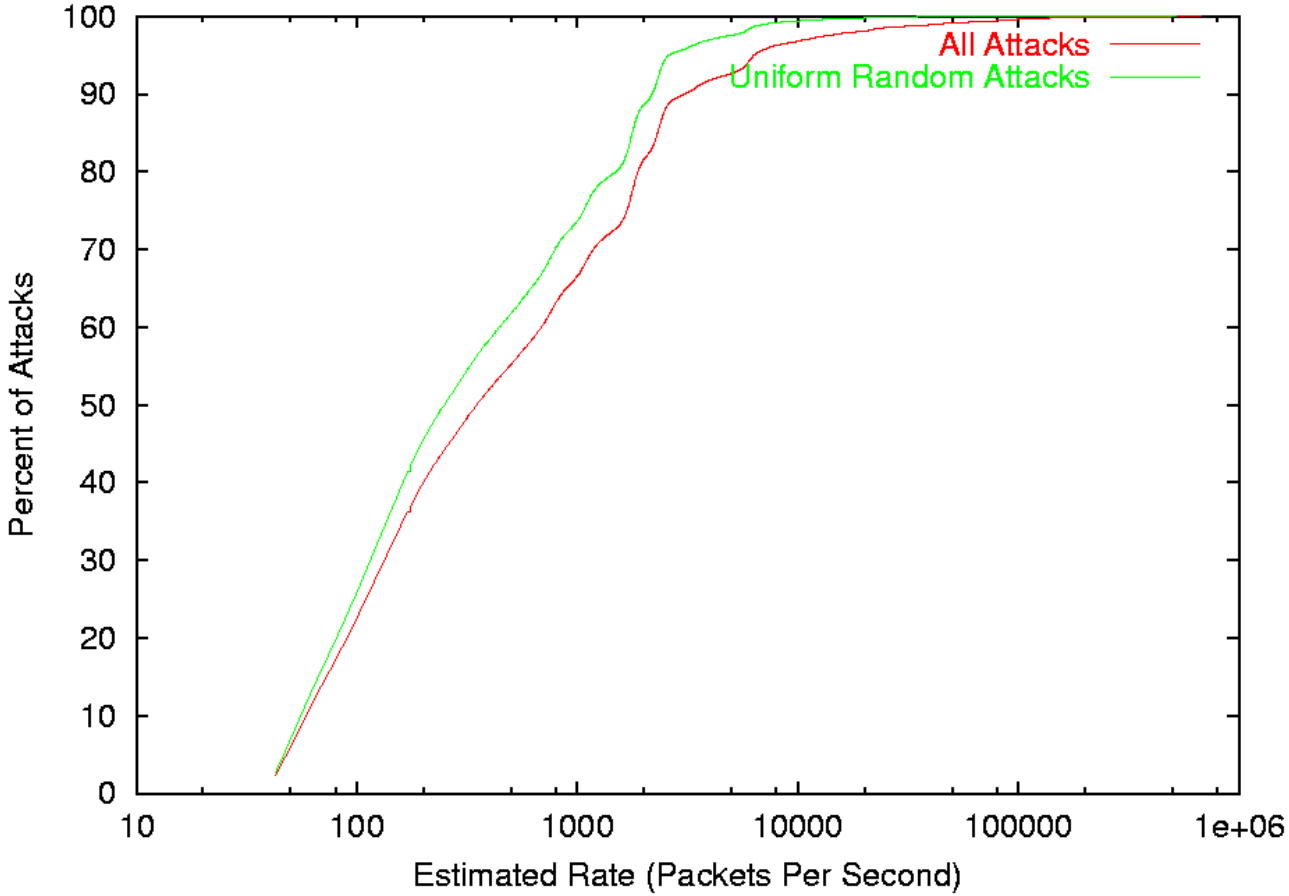

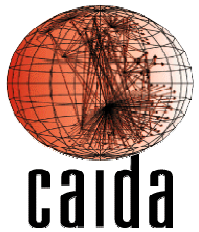

Victim characterization

- Entire spectrum of commercial businesses
  - Yahoo, CNN, Amazon, etc and many smaller biz
- Evidence that minor DoS attacks used for personal vendettas
  - 10-20% of attacks to home machines
  - A few very large attacks against broadband
- 5% of attacks target infrastructure
  - Routers (e.g. core2-core1-oc48.paol.above.net)
  - Name servers (e.g. ns4.reliablehosting.com)

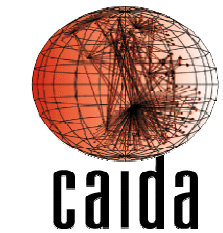

Victim breakdown by TLD

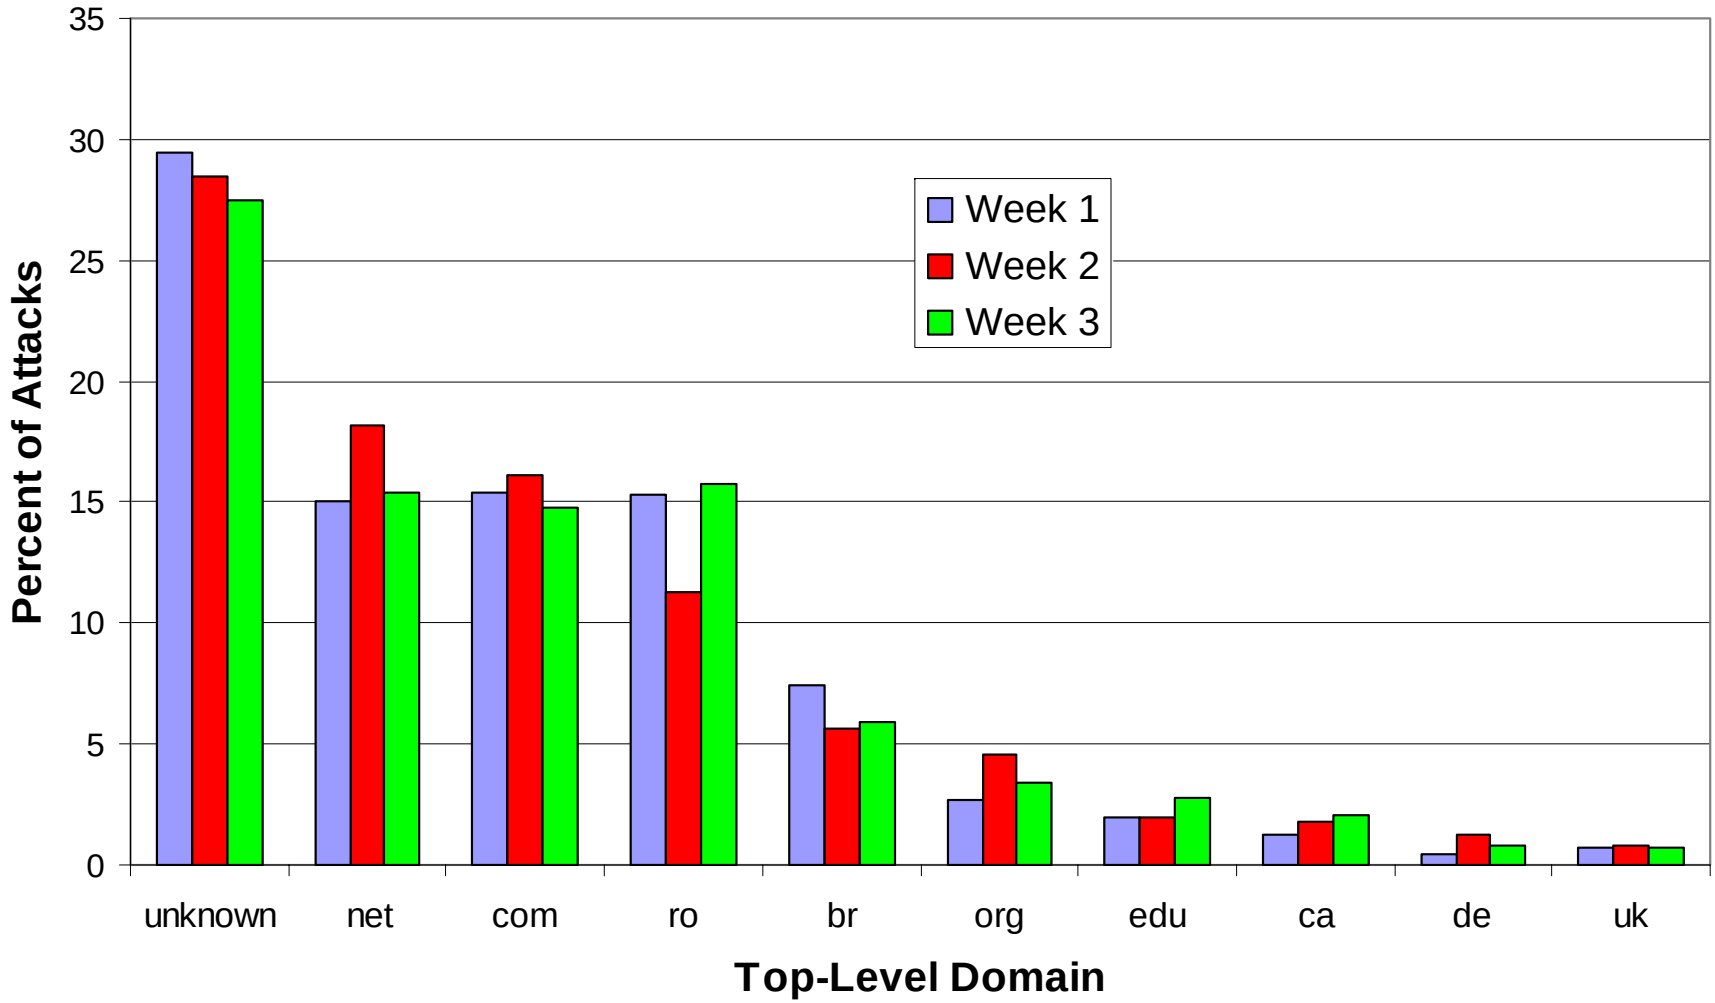

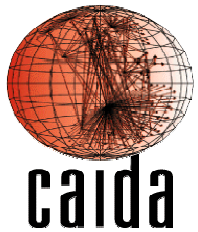

Distribution of repeat attacks

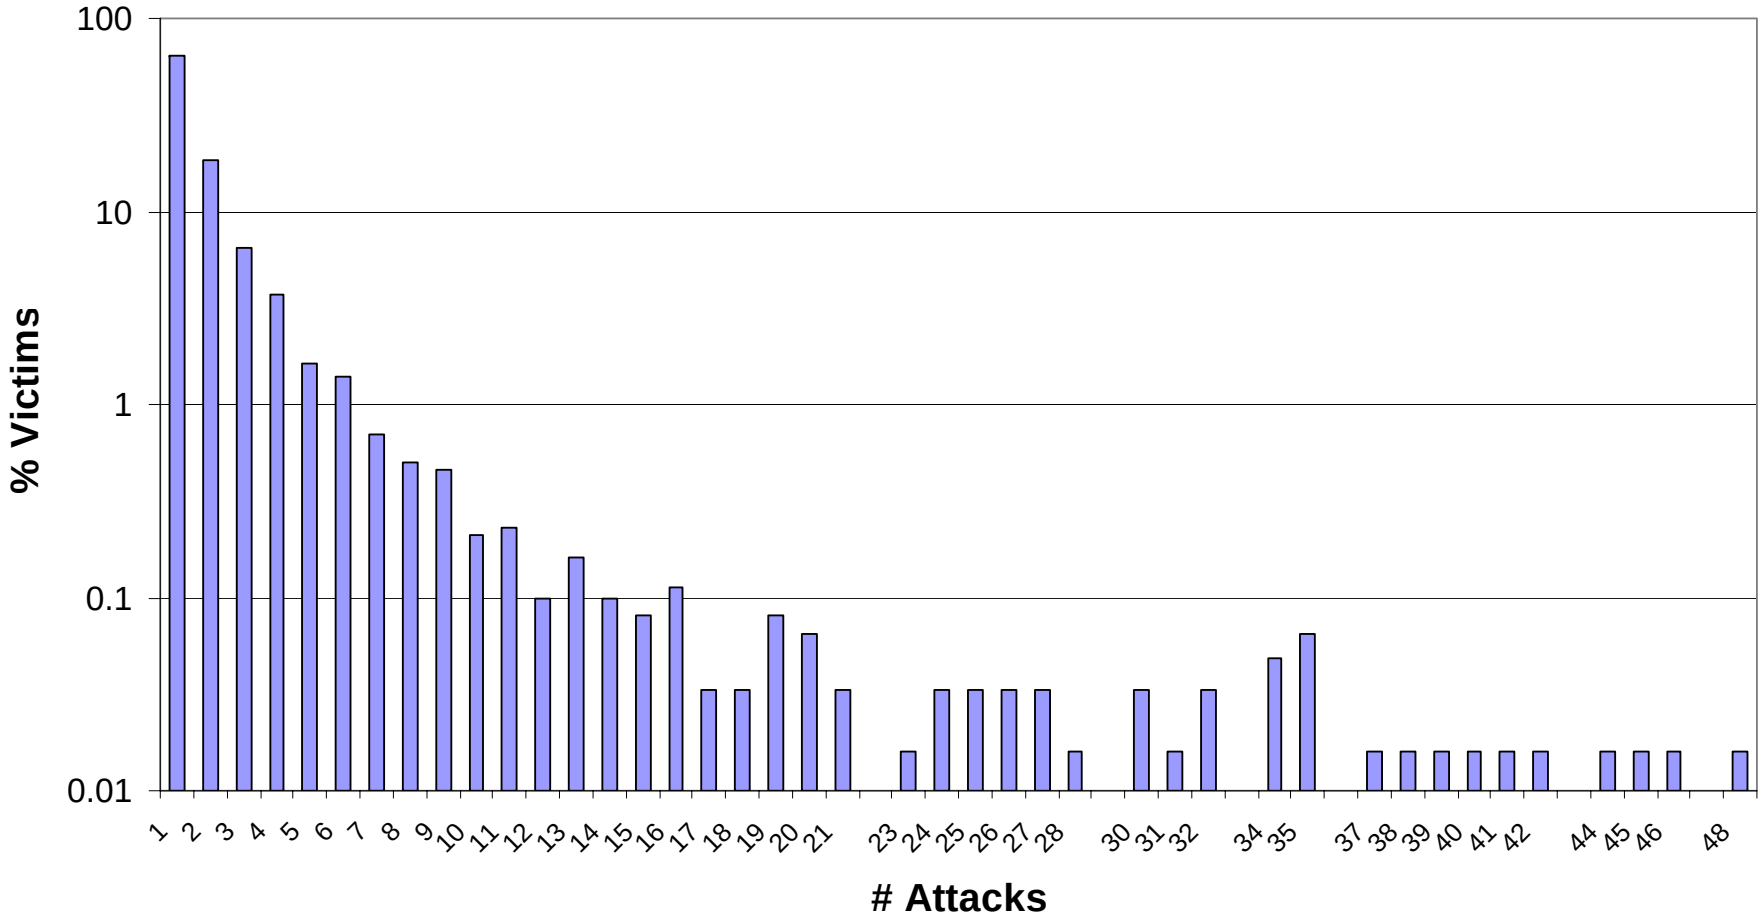

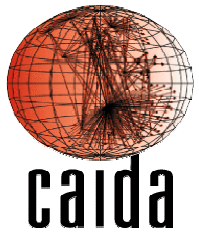

Validation

- Backscatter not explained by port scanning
  - 98% of backscatter packets don’t cause response
- Repeated experiment with independent monitor (3 /16’s from Vern Paxson)
  - Only captured TCP SYN/ACK backscatter
  - 98% inclusion into larger dataset
- Matched to actual attacks detected by Asta Networks on large backbone network

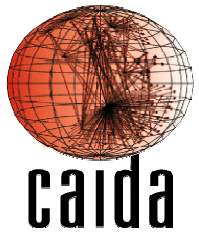

Conclusions

- Lots of attacks – some very large
  - **>12,000** attacks against **>5,000** targets
  - Most < **1,000** pps, but some over **600,000** pps
- Most attacks are short – some have long duration
  - a few victims were attacked continuously during the three week study
- Everyone is a potential target
  - Targets not dominated by any TLD, or domain
    - Targets include large e-commerce sites, mid-sized business, ISPs, government, universities and end-users
    - Targets include routers and domain name servers
  - Something weird is happening in Romania

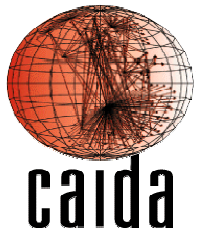

Cooperative Association for Internet Data Analysis
(CAIDA)
San Diego Supercomputer Center

Computer Science & Engineering
University of California, San Diego

http://www.caida.org/outreach/papers/backscatter/

Example 1: Periodic attack (1hr per 24hrs)

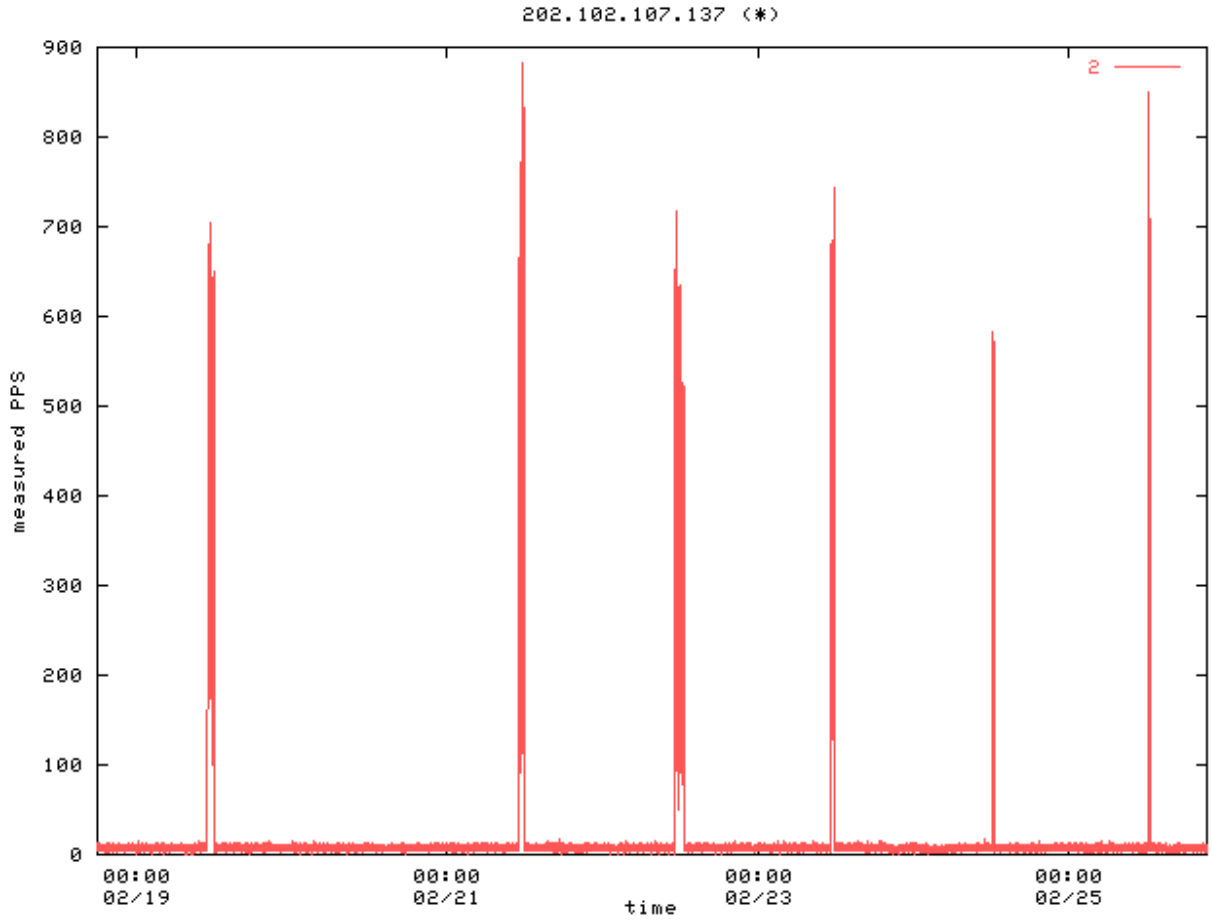

Example 2: Punctuated attack (1min interval)

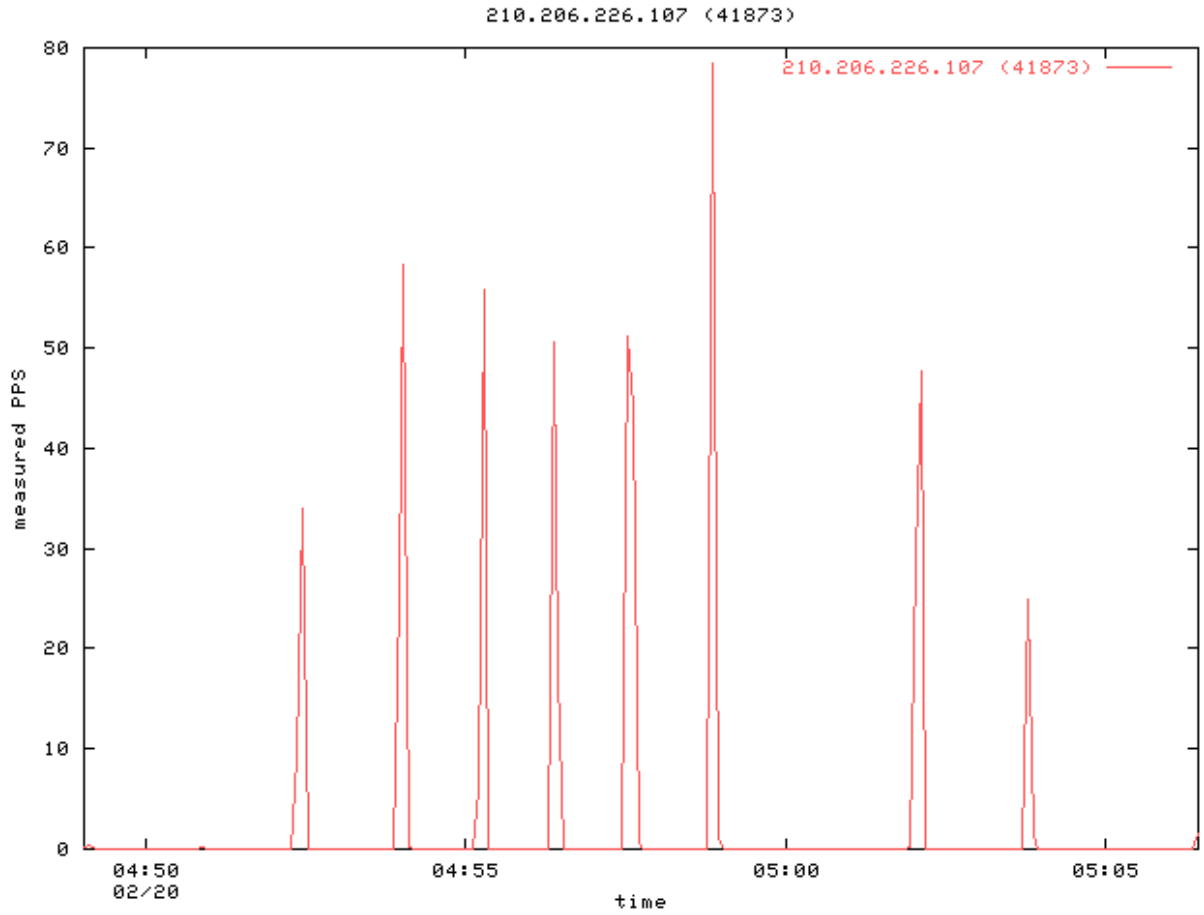

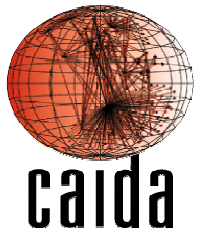

Backscatter protocol breakdown (one week)

Backscatter protocol	Attacks	BS Packets (x1000)
TCP (RST ACK)	2027 (49)	12,656 (25)
ICMP (Host Unreachable)	699 (17)	2892 (5.7)
ICMP (TTL Exceeded)	453 (11)	31468 (62)
ICMP (Other)	486 (12)	580 (1.1)
TCP (SYN ACK)	378 (9.1)	919 (1.8)
TCP (RST)	128 (3.1)	2,309 (4.5)
TCP (Other)	2 (0.05)	3 (0.01)

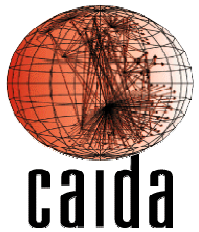

Attack protocol breakdown (one week)

Attack Protocol	Attacks	BS Packets (x1000)
TCP	3902 (**94**)	28705 (56)
UDP	99 (2.4)	66 (0.13)
ICMP	88 (2.1)	22,020 (**43**)
Proto 0	65 (1.6)	25 (0.05)
Other	19 (0.46)	12 (0.02)