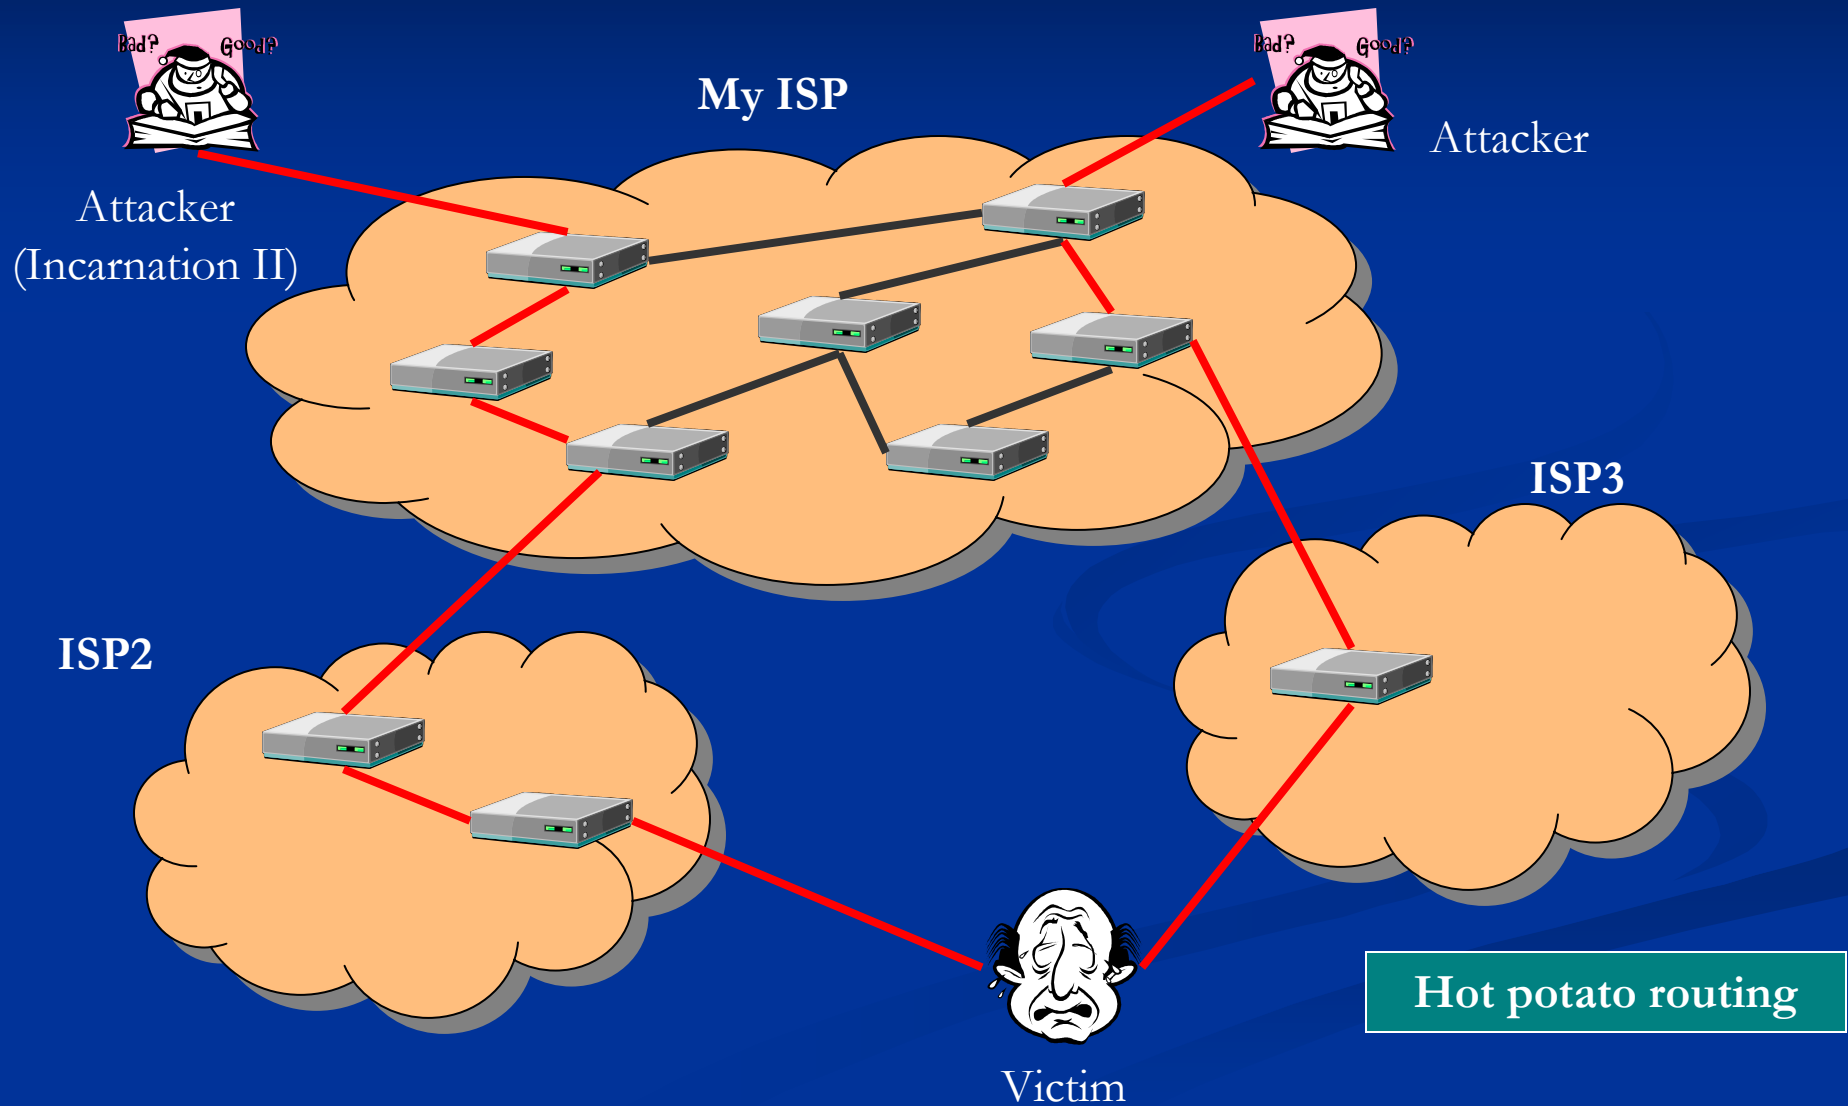


Detecting DDoS Attacks on ISP Networks

Ashwin Bharambe
Carnegie Mellon University

Joint work with:
Aditya Akella, Mike Reiter and
Srinivasan Seshan

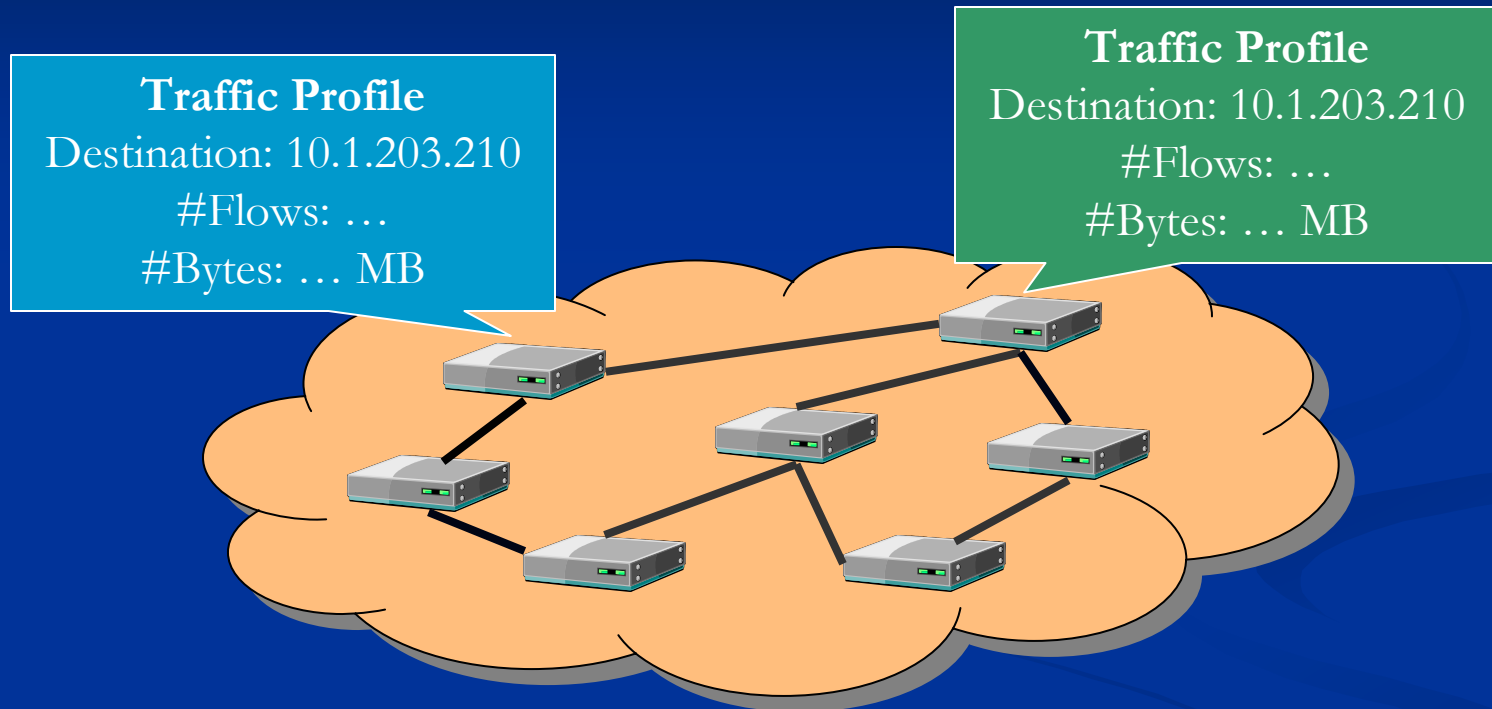
ISP Perspective of DDoS Attack



Problem Statement

- How can an ISP find out if:
 - Its Backbone is carrying “useless” attack traffic?
 - Its Backbone is itself under attack?
- Focus of this talk:
 - Sketch a solution approach
 - Discuss the main challenges

Approach



- Record “normal” traffic at routers; identify anomalies
- Exchange suspicious among routers to reinforce anomaly detection

Basic Approach

1. Record “normal” traffic at routers
2. Detect “abnormalities” in traffic

Challenges

- a. What is normal and what is abnormal?
- b. Is it robust?
- c. How quickly can we identify deviations?
- d. Can it really be implemented on a backbone router?
- e. Response strategy?

Proposed Solution

Maintain Traffic Profiles

- Each router constructs *profiles* of traffic
 - Longer time-windows → *normal* traffic
 - Smaller time-windows → *current* traffic
- **Become suspicious** if current profile **violates** normal profile

Important Challenges

1. Day-of-week and Time-of-day effects
 - Maintain per-day per-daytime statistics
2. Flash crowds
 - Example of “harmless” but infrequent event
 - Attack-volume alone is not a sufficient indicator
 - “Fingerprint” the destination-bound traffic
 - Number of sources, source-subnets, flows, distribution of flow lengths, etc.

Traffic Fingerprints

Some examples

- Total traffic to destination
- Source subnet characterization
 - Total number of “flows” to a destination
 - How many /24 subnets are observed in the traffic to this destination
- Flow-length distribution
 - E.g., are there a lot of small flows?

Stream Sampling

- Memory/computation constraints at routers
 - Keep statistics about every destination?
 - Only for popular ones \rightarrow traffic to whom exceeds a fraction θ of link capacity
 - Use **sample-and-hold** or **multistage filters** [Estan01]
 - Count unique subnets in a packet stream
 - Memory = $\Omega(\text{size of stream})!$
 - Use F_0 computation algorithms [Alon96, Gibbons01]
 - Do it in much smaller (constant!!) space and time

Proposed Solution

Increasing Robustness

- Single router has only local view → can make mistakes
 - Traffic perturbations due to traffic engineering
 - False alarms!
 - Suppose attacker “mimics” normal traffic at a router
 - Attack goes undetected!
- Mimicking at more than a few routers within an ISP would be hard!
- Use router consensus for reinforcing suspicions across routers

Preliminary Results

Single Router Detection Accuracy

Experimental Setup

- Abilene-II traffic trace (70 minutes)
 - Samples taken across a window of about 1 minute
- Synthetic attack traffic (trinoo, TFN, TFN2k, etc.)

Attack Detection Accuracy

- False positive rates $\leq 6\%$, lower for “unpopular” destinations
- False negative rates decrease rapidly as the “rate” of attack traffic increases

Conclusions and Future Work

■ Conclusions

- Fingerprinting traffic allows for detection of subtle attack patterns not apparent from volume alone
- Distributed detection makes it harder for an attacker to mimic traffic at multiple routers

■ Directions for future work

- Identify various attack scenarios
- Optimize computation/space requirements
- Consensus algorithm; convergence and effectiveness
- Validate over real attack datasets

Backup Slide

Overheads

Counting unique items in a stream (zeroeth moment F_0)

Algorithms	AMS96	GT01
Accuracy	$1 + \epsilon, \epsilon > 1$	$1 \pm \epsilon, \epsilon > 0$
Memory (bytes)	4	$36/\epsilon^2$
Byte operations	~ 4	~ 6

- Use $\epsilon = 0.1 \rightarrow$ memory ~ 3600 bytes *per destination*
- Approximate number of popular destinations = $1/\theta$
where θ is the fraction of link capacity
- 360 KB per statistic – if we use $\theta = 1\%$
- Can a high-end router have a few MBs of SRAM?