Adaptive Defense Against Various Network Attacks

Cliff C. Zou, Univ. Central Florida
Nick Duffield, AT&T Labs Research
Don Towsley, Weibo Gong, Univ. Massachusetts at Amherst

- Motivation and big picture
- System design #1 - SYN flood DDoS attack
- System design #2 - Internet worm attack
- Summary and future work

Outline

- Motivation and big picture
- System design #1 - SYN flood DDoS attack
- System design #2 - Internet worm attack
- Summary and future work

Motivation of Adaptive Defense

P_p : False positive prob.
blocking normal traffic

P_n : False negative prob.
missing attack traffic

θ : Detection sensitivity

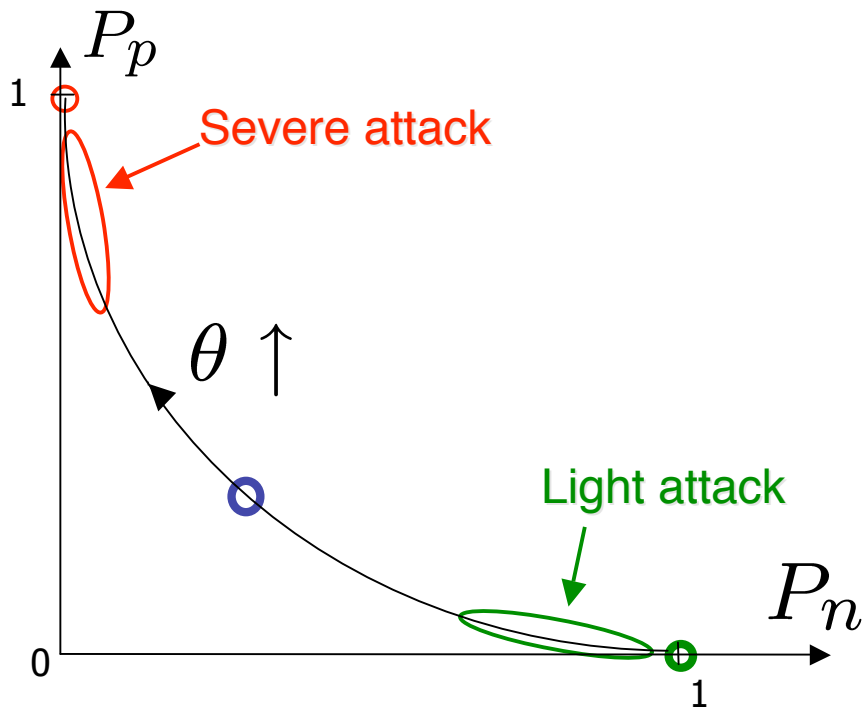

Q: Which operation point is "good"?

A: All operation points are good
Optimal one depends on attack severity

Adaptive Defense Principle

- More severe attack, more aggressive defense (with more false alarm cost)
  - Comparing with attack damage, we are willing to pay certain false alarm cost
  - Used in epidemic control in the real world
  - Implementation:

    Min (false alarm cost + missed attack cost)

Outline

- Motivation and big picture
- System design #1 - SYN flood DDoS attack
- System design #2 - Internet worm attack
- Summary and future work

SYN flood DDoS attack

- Attack hosts send TCP connection requests faster than a server can process them
  - Mostly with spoofed source IPs
- Filtering defense
  - Must based on individual TCP/SYN packet
  - Hop-Count Filtering – packet's TTL value [CCS'03]
    - Attackers don't know hop-counts from real clients to a server
    - It is the underlying detection algorithm we use

Estimation of attack severity π

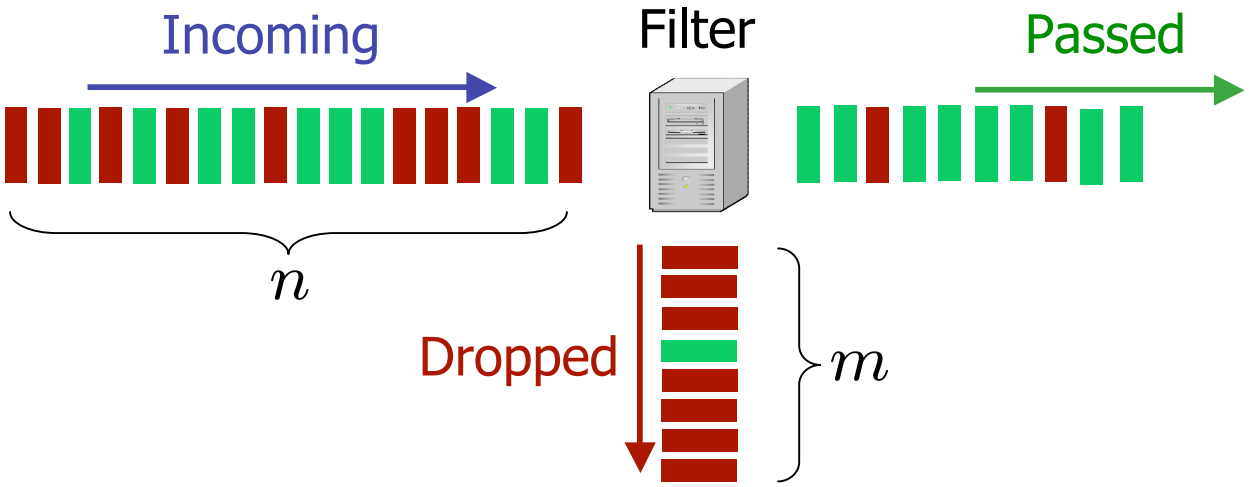

$\pi' \equiv \frac{m}{n}$: Fraction of detected attack traffic

$$m = \pi' n$$

$$\Rightarrow \quad \hat{\pi} = \frac{\pi' - P_p}{1 - P_n - P_p}$$

(π': # of incoming attack traffic; # of incoming traffic)

$$E[\hat{\pi}] = \pi \quad \text{Unbiased}$$

Adaptive Defense Design

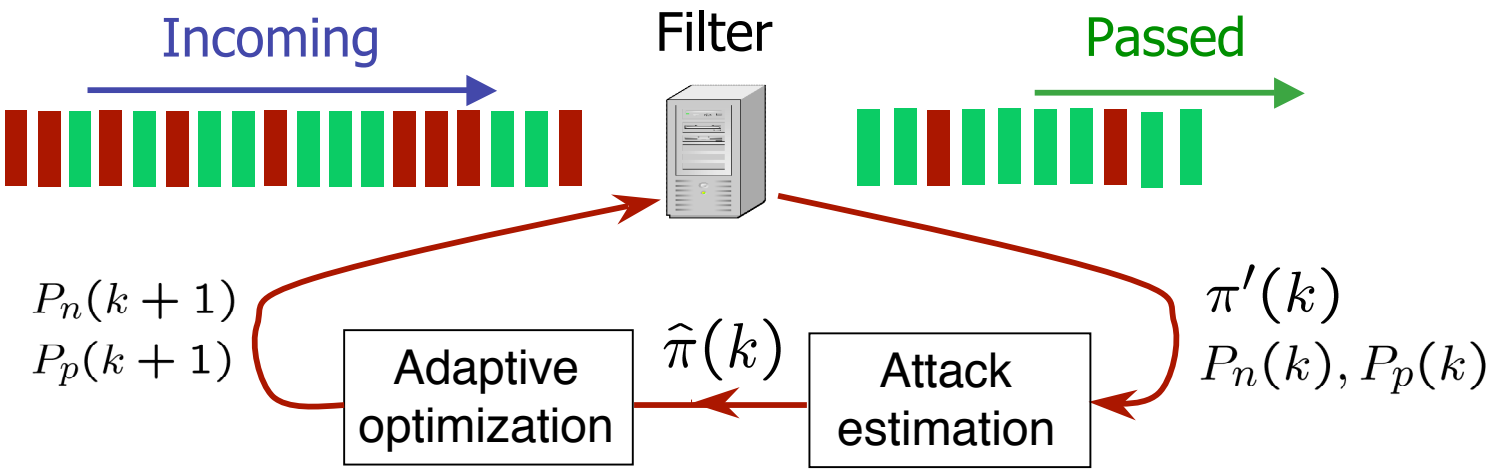

Discrete time k → k+1

Optimization:

c_p : Cost of dropping a normal traffic (Fraction of dropped normal)

c_n : Cost of passing an attack traffic (Fraction of passed attack)

Adaptive Defense Results

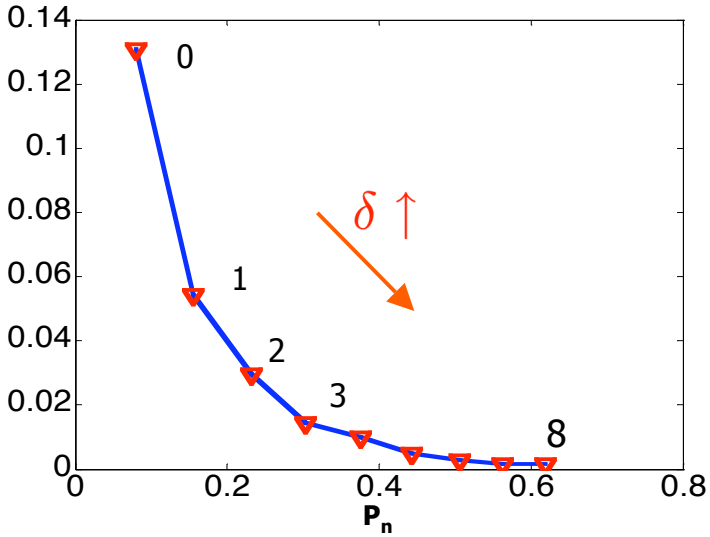

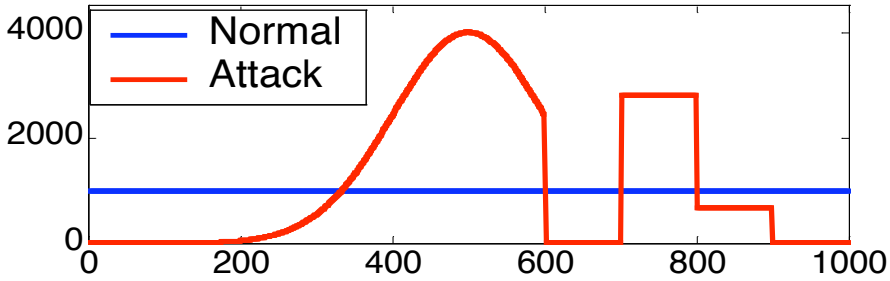

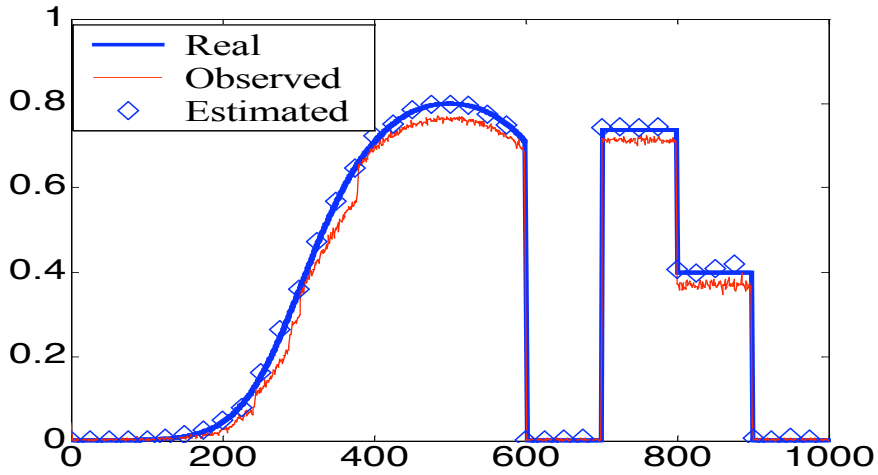

Adaptive Defense Results

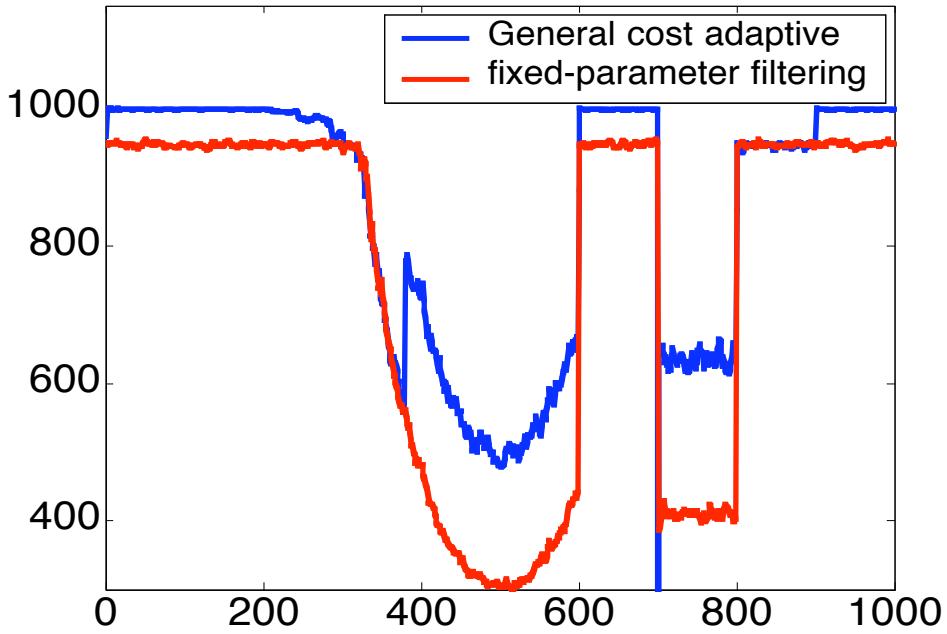

- Adaptive defense is better when
  - Under normal situation
  - Under severe attacks

Outline

- Motivation and big picture
- System design #1 - SYN flood DDoS attack
- System design #2 - Internet worm attack
- Summary and future work

Internet Worm Attack

- Protect a local network from outside worm infection
- Local detection (without signature)
  - Modified Threshold Random Walk [IEEE S&P'04]
    - failed connections >> success connections
- Defense : Black-listing on edge routers
  - TCP worms
  - UDP worms without spoofing

Adaptive Defense Design

- Modified Threshold Random Walk ([Usenix'04])
  - Receive a failed request → the source's counter + 1
  - A success request → the source's counter – 1 (if >0)
  - Counter ≥ W → Mark the source as an attacker

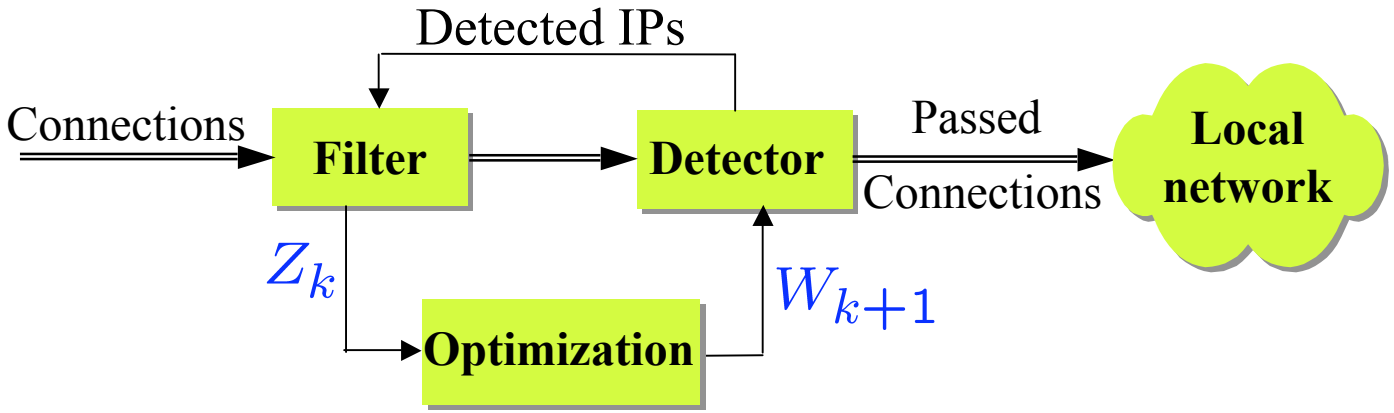

Adaptive Defense Results

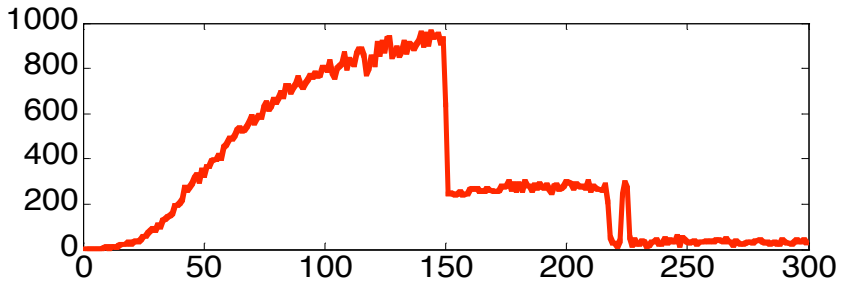

W

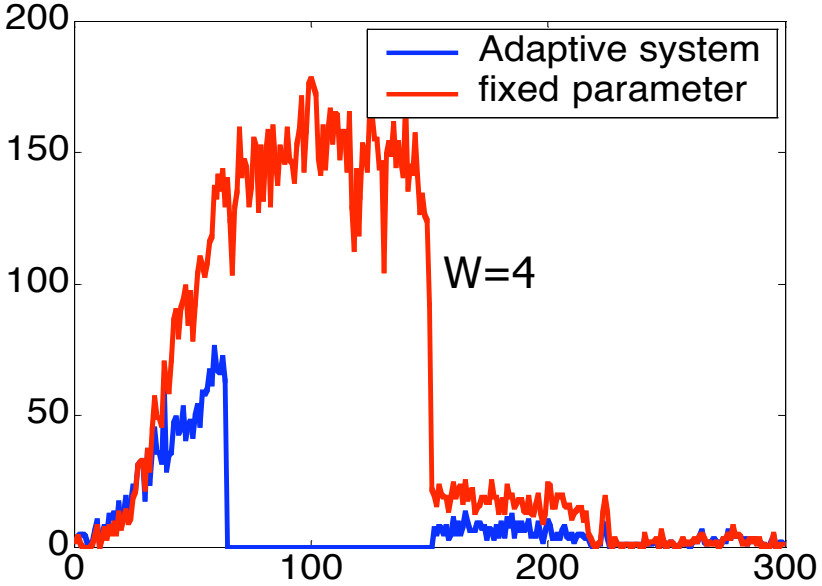

- Slammer monitored trace (from Andrew Daviel)
  - /16 network monitoring
  - Observed nearly 10,000 attack sources in the first 5 minutes.

Outline

- Motivation and big picture
- System design #1 - SYN flood DDoS attack
- System design #2 - Internet worm attack
- Summary and future work

Adaptive Defense Summary

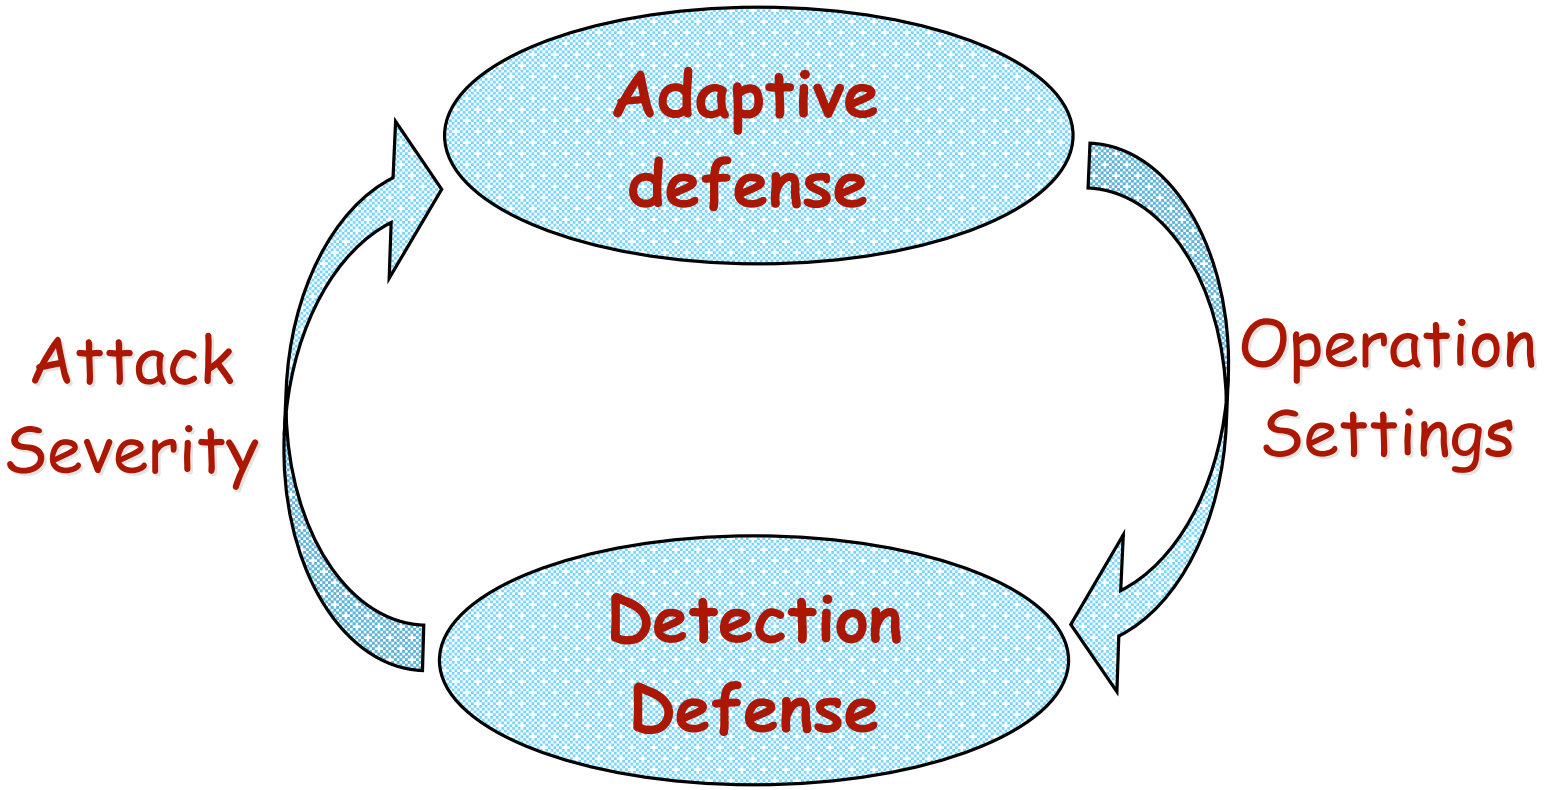

More severe attack, more aggressive defense

Future Work

- System evaluation:
  - Real trace with both normal and attack traffic
  - On more underlying detection algorithms
- How to determine penalty factors c_p, c_n ?
- How to define **cost** when:
  - P_p, P_n are not clearly defined?
  - Detection time is critical?
- Tunable by attackers?
  - Cautious in using attack prediction