Optimal Allocation of Filters against DDoS Attacks

Karim El Defrawy, Athina Markopoulou
UC Irvine
Katerina Argyraki
EPFL, Lausanne

The general DDoS problem

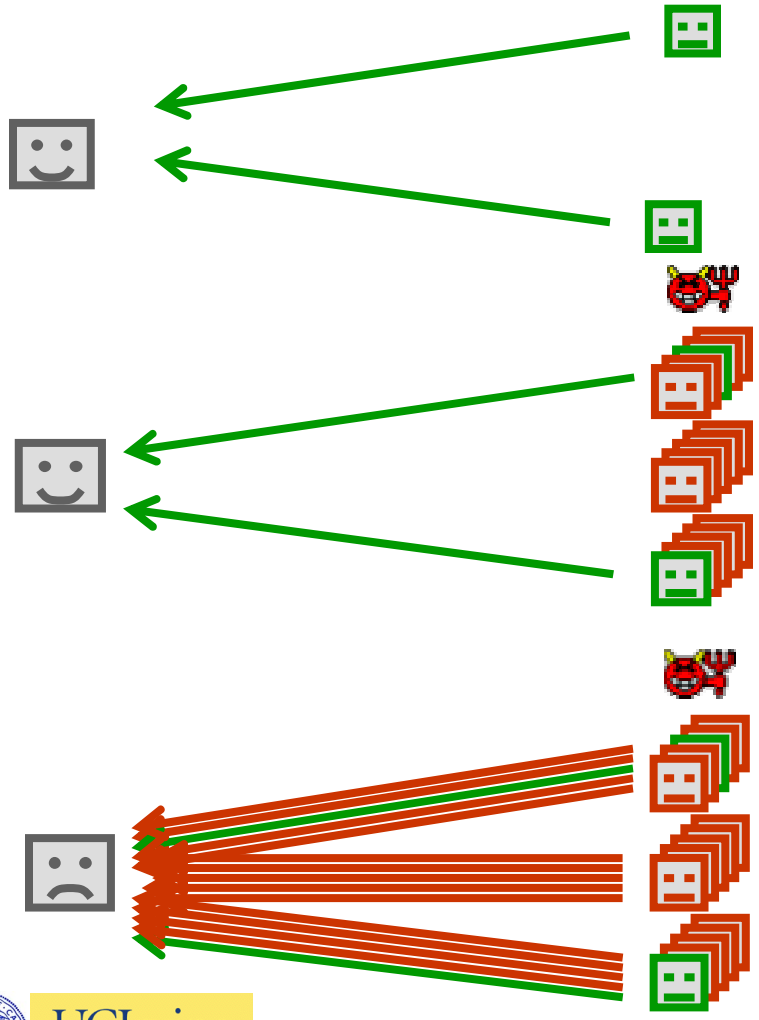

- Good users can access the server and have a good experience
- Some hosts are compromised. No attack is launched yet.
- The attack is launched. The server has no resources left for the good users.

Distributed Denial-of-Service (DDoS) Attacks

- What is DDoS?
  - A large number of attacking nodes (tens of thousands) send huge amount of traffic consuming the victim's resources, such as CPU, memory, bandwidth.

- A severe problem in the Internet today
  - Akamai, DoubleClick, Amazon, disrupted for hours
  - eBay: 1 hour of downtime = $180,000
  - bluesecurity.com
  - www.whitehouse.gov, www.cnn.com, …
  - High frequency of attacks (1000s per week) reported

The Flooding Attack
(attack on the tail-circuit bandwidth)

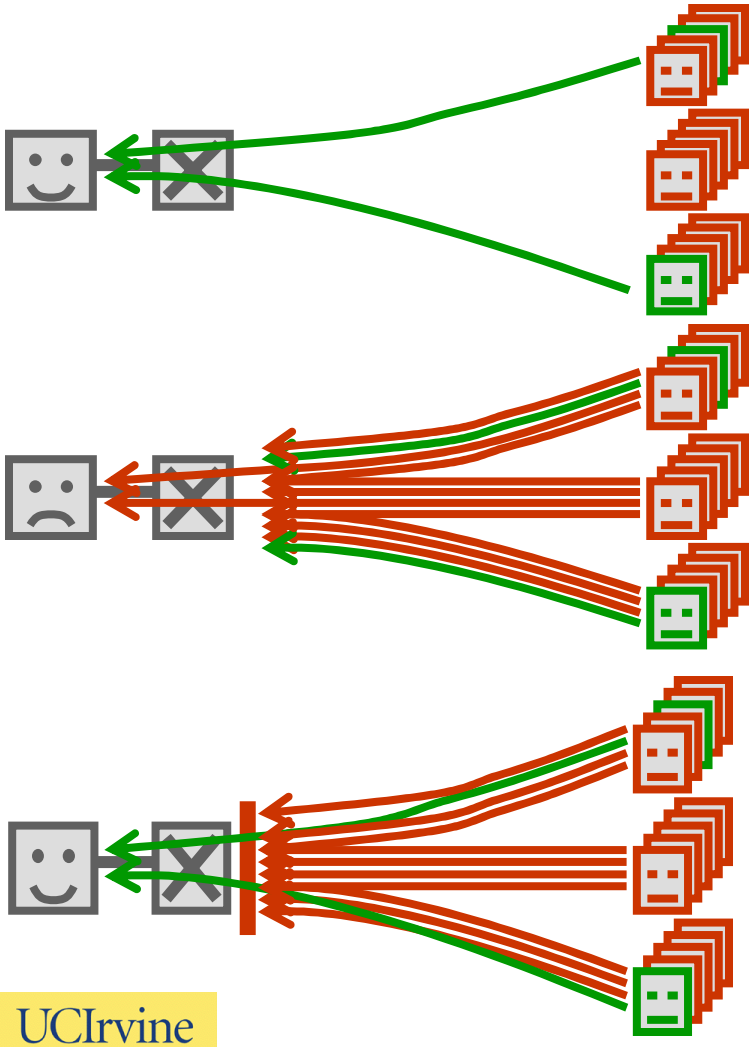

- Good users access the server.
- Some hosts are compromised.

- Flooding attack launched. → tail-circuit bandwidth exhausted.
- Good users back off, goodput →0
- Inherent weakness of the Internet paradigm.

- Several solutions proposed. We focus on filtering.

One Defense Mechanism: Filtering

- What are Filters?
  - Access Control Lists (ACLs) can match a packet header against rules, e.g. source and destination IP addresses.
  - Rate limiters
- Filters are an expensive resource
  - stored in TCAM
    - 1 TCAM chip per router linecard
    - at most 256K filters per TCAM chip
    - each victim gets only a few 1000s of filters
- There are more attackers than filters
  - An attack can consist of millions of flows

Problem Statement

- Goal:
  - Allocate filters to attackers or groups of attackers (at a single router) so as to minimize the damage caused by the DDoS attack, subject to constraints in the #filters and the capacity ?

- Contributions:
  - Optimization problems
  - Demonstrate that optimal allocation has significant benefit in practical cases
  - Develop efficient heuristics

The Filter Allocation Problem

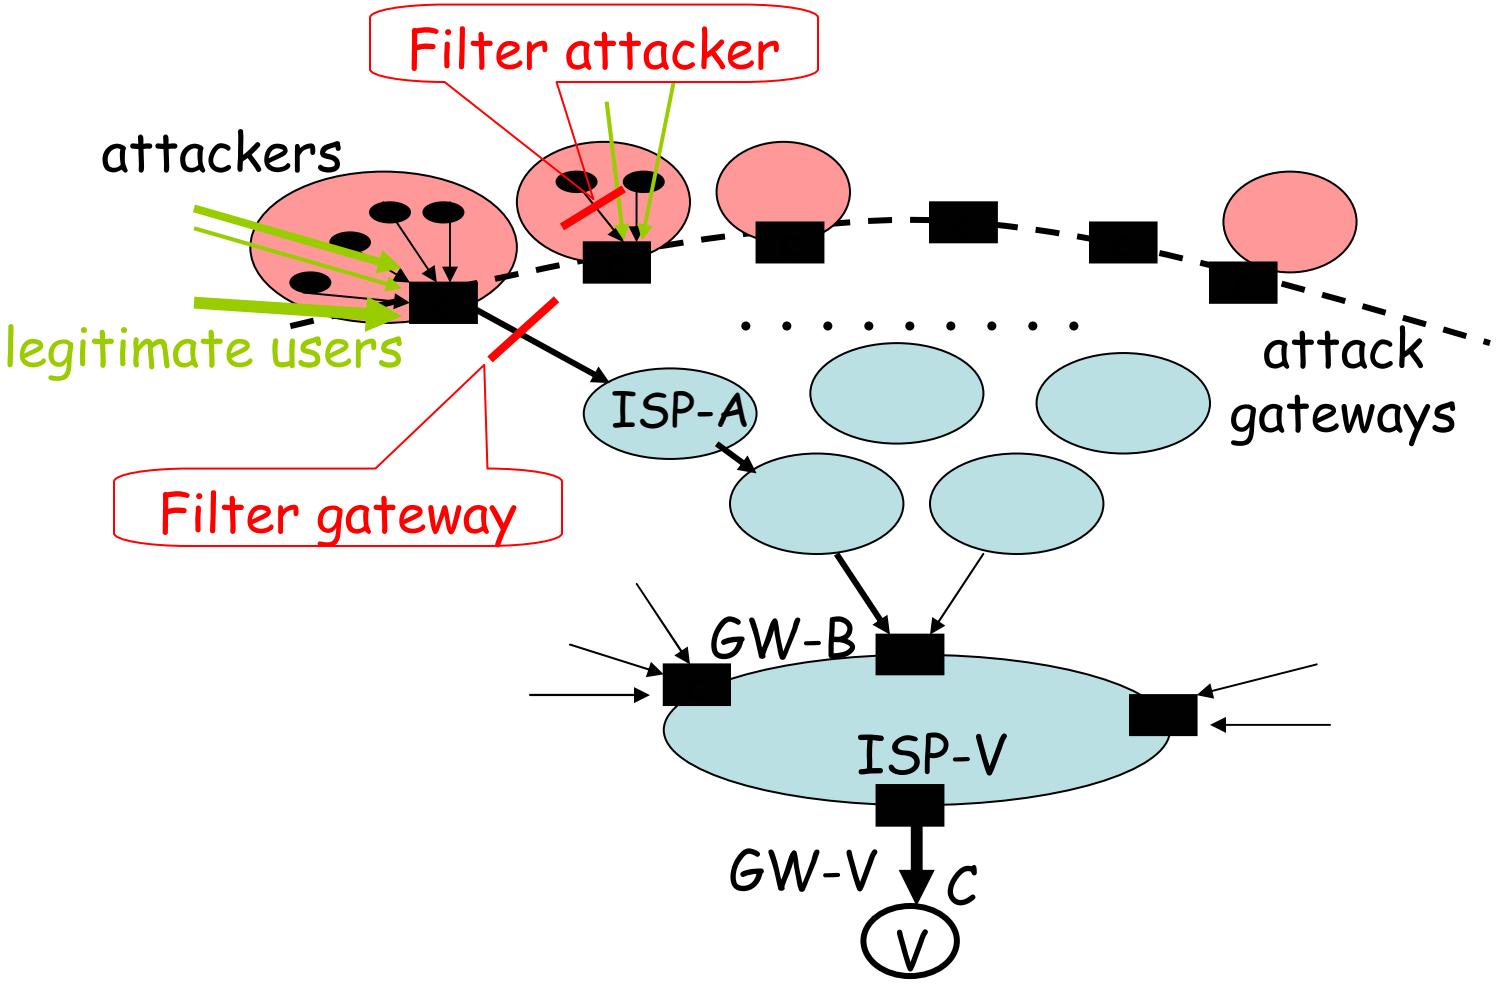

Tradeoff: saving filters vs. collateral damage

Single-Tier Filter Allocation
(at gateway-level only)

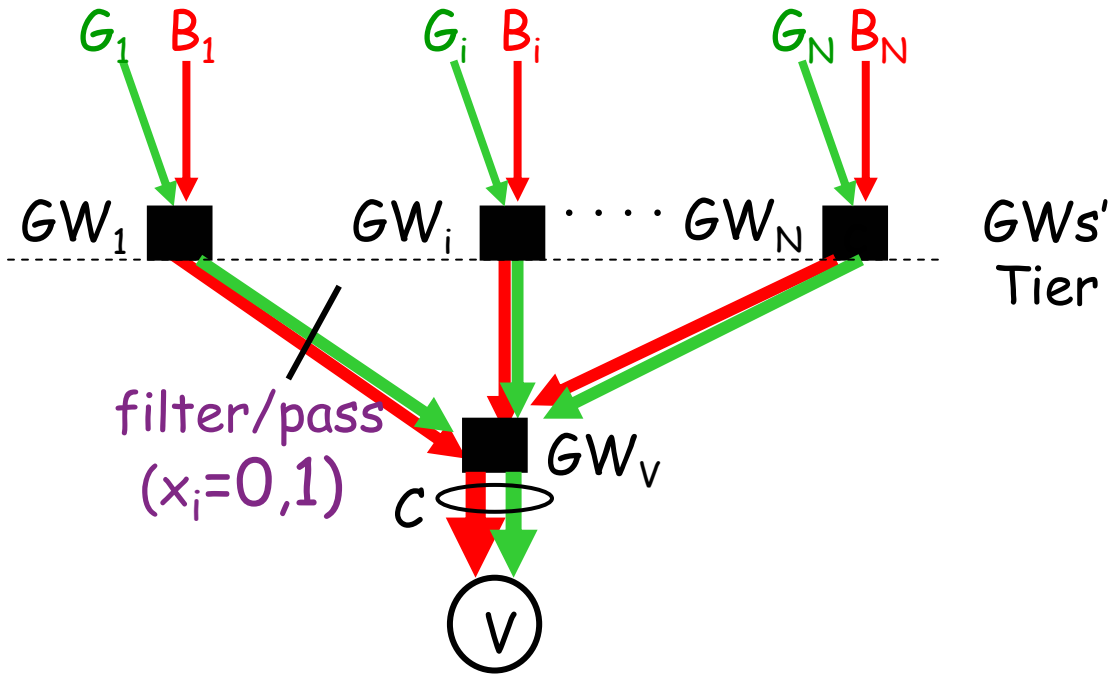

$$\max \sum_{i=1}^{N} G_i x_i$$

$$\text{s.t.} \sum_{i=1}^{N} (G_i + B_i) x_i \leq C$$

$$\text{and } x_i = 0,1,\ i = 1,2,\ldots N$$

UCIrvine
University of California, Irvine

Single-Tier Optimal Solution

- A 0/1 knapsack with N items.
  - item i has profit G_i and cost (G_i+B_i)
- The fractional problem $(0 \le x_i \le 1)$
  - solved by a greedy alg: order in decreasing $G_i/(G_i+B_i)$
  - has optimal solution ($x_1=1,\ldots x_{c-1}=1$, $x_c<=1$, $x_{c+1}0,\ldots x_N=0$)

pass — only 1 rate limiter — filter

$$\max \sum_{i=1}^{N} G_i x_i$$

$$\text{s.t.} \sum_{i=1}^{N} (G_i + B_i) x_i \leqslant C$$

$$\text{and } x_i = 0,1,\ i = 1,2,\ldots N$$

Two-Tier Filter Allocation
general version

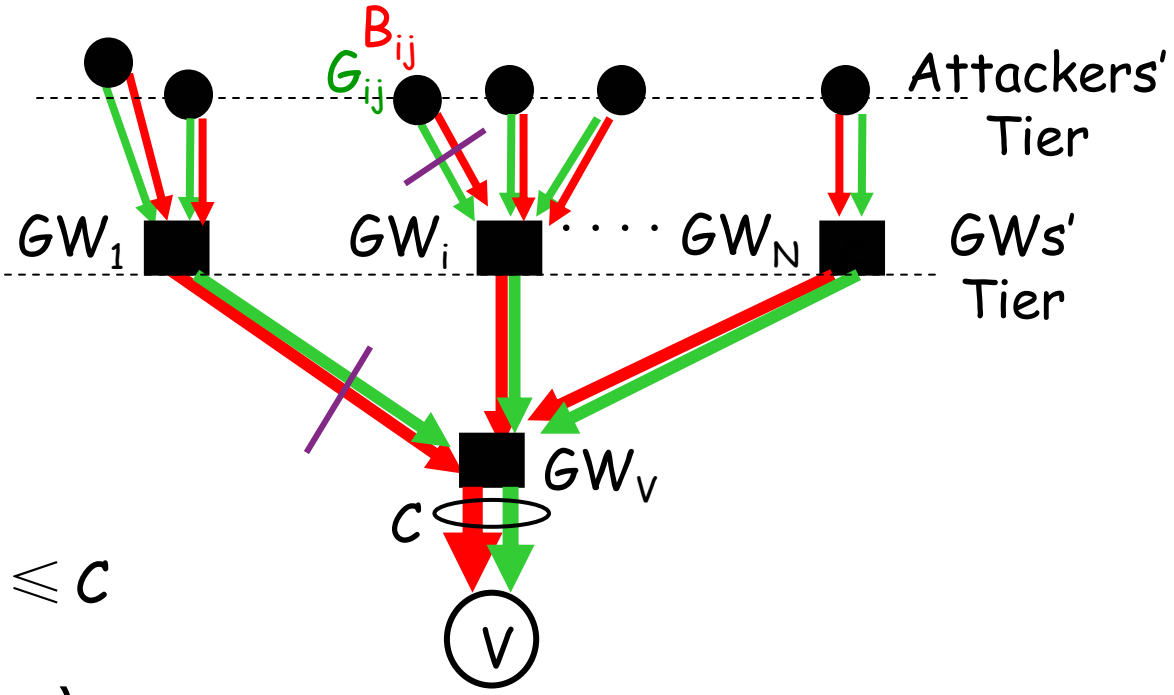

$$\max \sum_{i=1}^{N} G_{ij} x_i x_{ij}$$

$$\text{s.t.} \sum_{i=1}^{N} \sum_{j=1}^{M_i} (G_{ij} + B_{ij}) x_i x_{ij} \leq C$$

$$\sum_{i=1}^{N} (1 - x_i) + \sum_{i=1}^{N} \sum_{j=1}^{M_i} (1 - x_{ij}) \leq F$$

$$x_i, x_{ij} \in \{0,1\}, \; i = 1,2,\ldots N, \; j = 1,\ldots M_i$$

attacker-ij filter/pass
GW-i filter/pass

Two-Tier Filter Allocation

assuming separate hosts for attackers and users

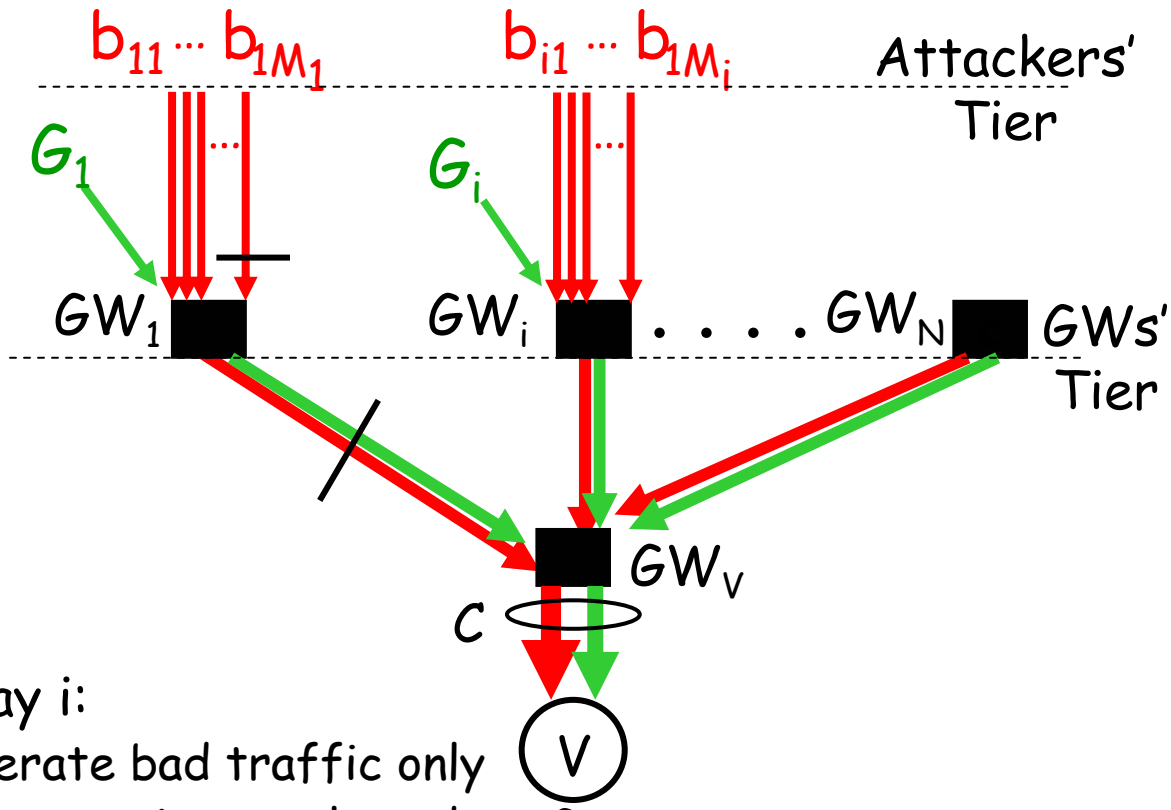

- Behind each gateway i:
  - Attack hosts generate bad traffic only
  - Different hosts generating total goodput G_i
  - Total traffic per gateway C_i
- Maximize goodput $T_N(C,F)$

UCIrvine
University of California, Irvine

Optimal Two-Tier Allocation
Dynamic Programming Formulation

- Let $T_n(c,f)$ be
  - the maximum goodput, considering gateways {1,2,..n}, using f<=F filters and capacity c<=C
- We can calculate $T_N(C,F)$ iteratively.
  - by filling up a table $T_n(0{:}C,0{:}F)$ for n=1,2,...N
- In iteration n:
  - consider all gateways up to n: {1,2...n}. Two groups: {1,2...n-1} and {n}
  - consider assigning (f-x) filters to {1,...n-1} and x filters to {n}
  - choose the best number 0<=x<=f and configuration

$$T_n(c,f) = \max_{x=0,1,...f}\{T_{n-1}\left(c-(C_n-\sum_{j=1}^{j=x}b(n,j)),f-x\right)+G_n\}$$

- x=0: $T_{n-1}(c-C_n,f)+G_n$
- x=1: $\max\{T_{n-1}(c,f-1),T_{n-1}(c-(C_n-b(n,1)),f-1)+G_n\}$

Simulation Setup

- Filtering performance
  - in terms of %goodput preserved, #filters used
  - depends on the distribution of good/bad traffic
- Compare optimal filtering to other policies:
  - No filtering, random filtering, uniform rate limiting, max-min
- Distribution of bad traffic:
  - Code Red [CAIDA], Slammer, Zombie Report [Prolexic]
    - They all provide the % of infected hosts per AS.
  - Uniformly spread attack scenario
- Distribution of good traffic:
  - % Internet users per country [InternetWorld stats]
- Other parameters
  - Consider C=100Mbps, all flows send at DSL rates
  - Vary the number of bad and/or good flows.

Attack Scenario: Code Red

		Code Red I		Code Red II	
Country	GW	% of Good Traffic from [20]	% of Bad Traffic from [16]	% of Good Traffic	% of Bad Traffic
USA	1	36.27	43.9	36.2	45.9
Korea	2	5.8	11.5	0	12
China	3	18.35	10.3	24.1	0
Taiwan	4	2.46	6.1	2.4	16.7
Canada	5	3.64	5.4	3.6	5.4
UK	6	6.74	5.2	6.7	5.3
Germany	7	8.4	5.1	8.4	5.2
Australia	8	2.5	4.3	2.5	1.1
Japan	9	13.91	4.2	14.2	0
Netherlands	10	1.93	4.1	1.9	8.4

UCIrvine
University of California, Irvine

Attack Scenario: Slammer

Country	GW	% Good Traffic	% Bad Traffic
USA	1	36.3%	44.6%
South Korea	2	5.8%	13.6%
China	3	18.5%	8%
Taiwan	4	2.4%	5.7%
Canada	5	3.6%	4.6%
Australia	6	2.5%	4.2%
UK	7	6.7%	3.8%
Japan	8	13.9%	3.5%
Netherlands	9	1.9 %	3.3%
Unknown	10	8.4%	8.7%
Total		100%	100%

Attack Scenario: Zombie Report [Prolexic]

Country	GW	% Good Traffic	% Bad Traffic
US	1	36.5%	21.5%
China	2	18.5%	14.5%
Germany	3	8.5 %	13.5%
UK	4	6.78%	8.5%
France	5	4.59%	8.5%
Brazil	6	4%	7.5%
Japan	7	13.99%	7.5%
Phillippines	8	1.4%	6.5%
Russia	9	13.94%	6.5%
Malaysia	10	1.8%	5.5%
Total		100%	100%

Uniformly spread attack scenario

- Same number of hosts behind each GW
- N good hosts are chosen at random
- M bad hosts are chosen at random
- All hosts emit at the same rate (DSL)
- Mixed:
  - Nodes can be both good and bad (emit both good and bad traffic)
- Not Mixed
  - Nodes can be either good or bad (emit only good or only bad traffic

Optimal Two-Tier Filtering
against a Code-Red Based Attack

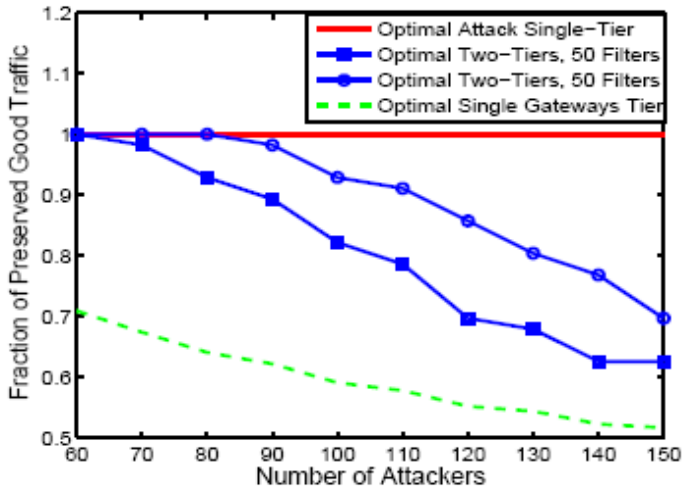

(a) % Good Traffic Preserved

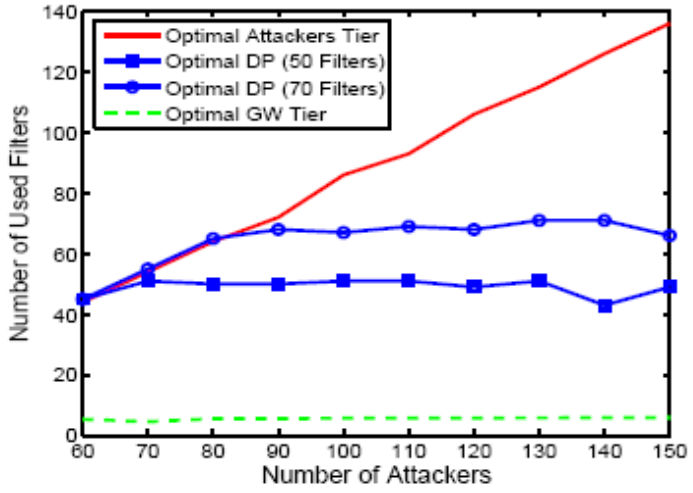

UCIrvine
University of California, Irvine

Optimal Two-Tier Filtering
against a Slammer Based Attack

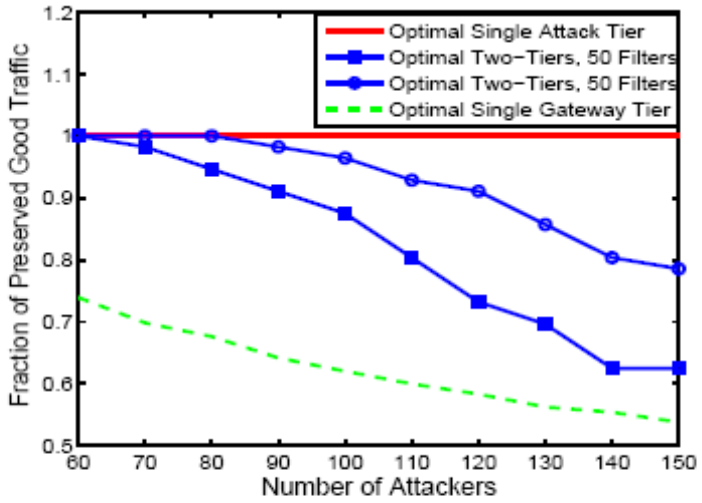

(a) % Good Traffic Preserved

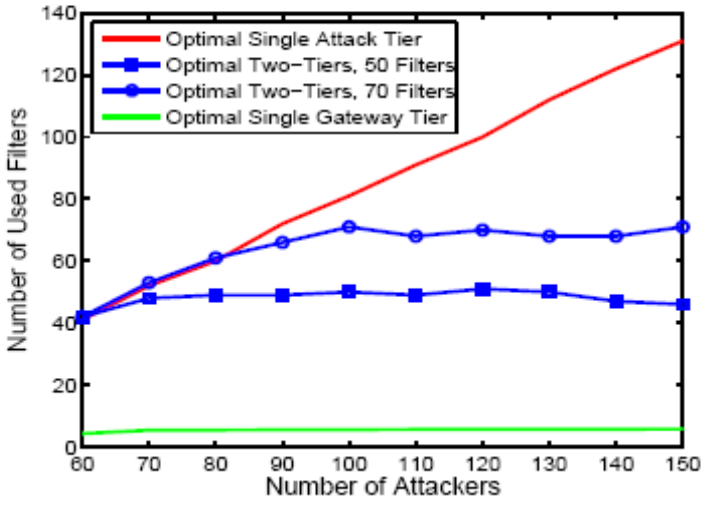

UCIrvine
University of California, Irvine

Two-Tier Optimal Filtering
against a Zombie Report Based Attack

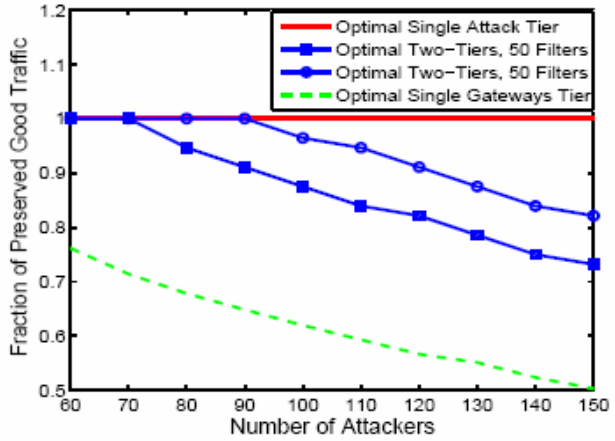

(a) % Good Traffic Preserved

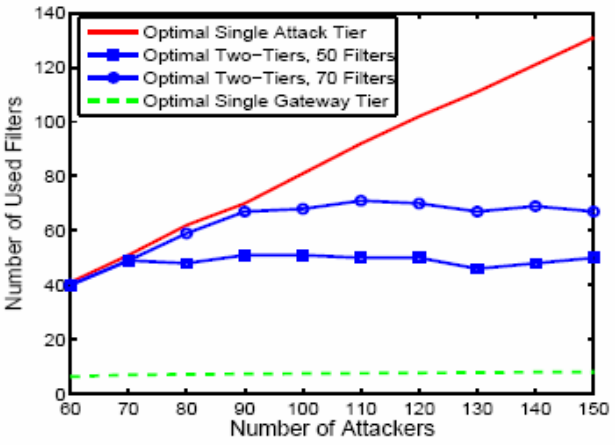

(b) Number of filters used.

UCIrvine
University of California, Irvine

Optimal Single (GW) Tier Filtering
against a Code-Red Based Attack

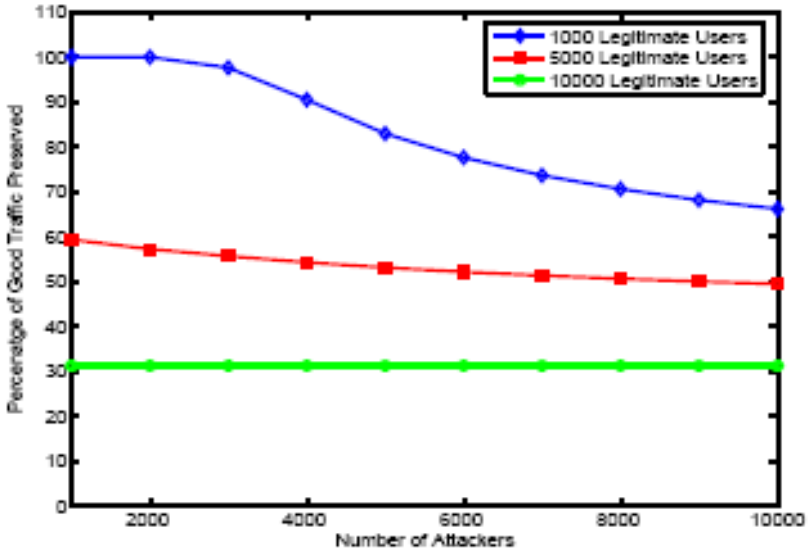

(a) % Good Traffic Preserved after Optimal Filtering

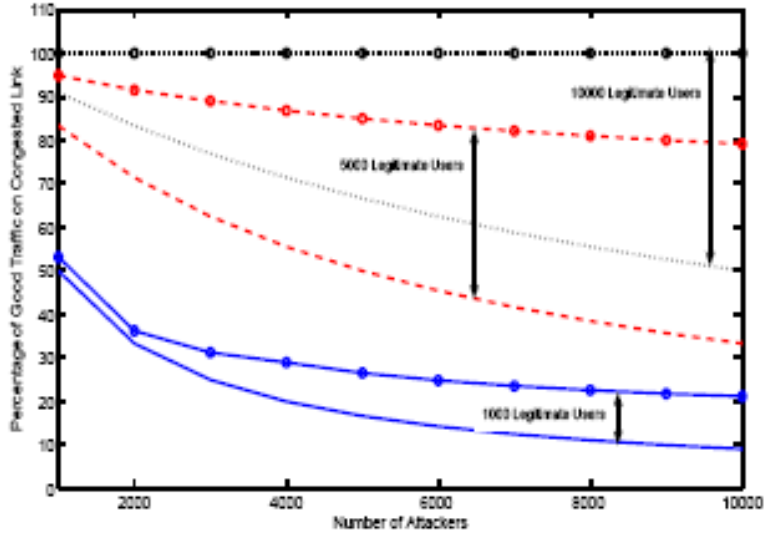

(b) Good % Link BW before/after Optimal Filtering.

Optimal Filtering preserves all good traffic
(unless it exceeds the capacity)

UCIrvine
University of California, Irvine

Single Tier: Attackers' vs. GWs' Tier

uniform, mixed attack scenario

Attackers' single tier

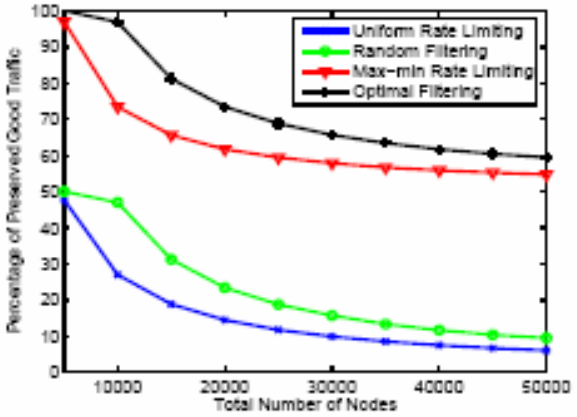

(a) % Good Traffic Preserved after Filtering

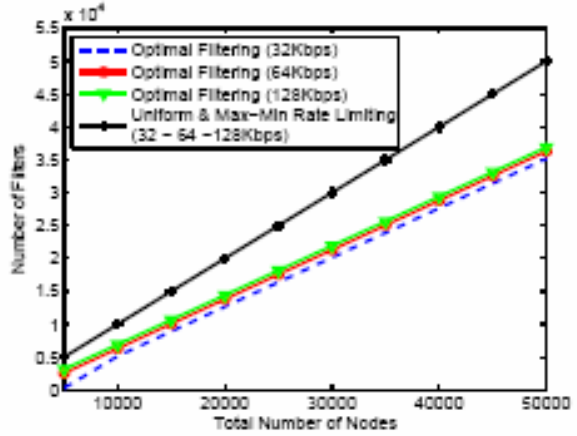

(b) Number of filters used.

Gateways' single tier

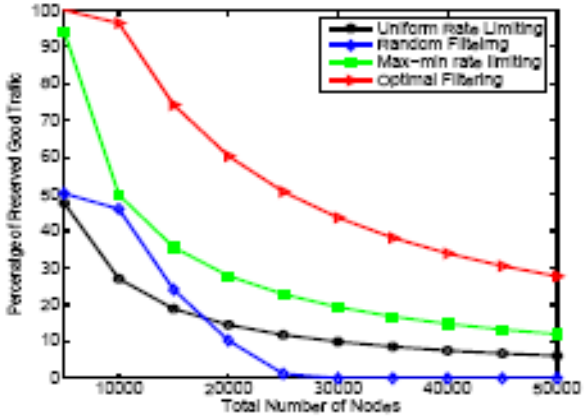

(a) % Good Traffic Preserved after Filtering

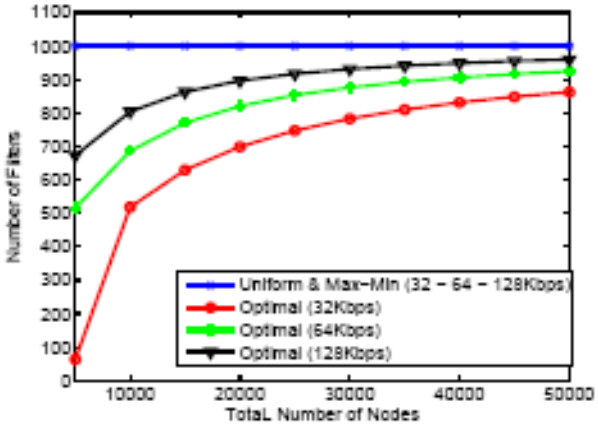

(b) Number of filters used.

Need for Heuristics

- We need low-complexity so it can
  - Compute and respond real-time
  - Re-compute to adjust to dynamic attacks

- Single-tier optimal solution fast but coarse
- Multi-tier optimal solution takes O(NCF)

- Need for low-complexity algorithms

Heuristics: attacker-based vs. GW-based

for the mixed uniform scenario

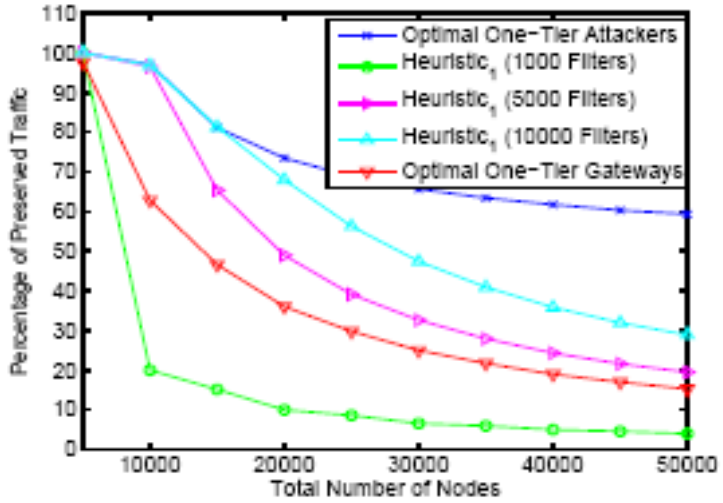

(a) Heuristic 1(Attacker-based)

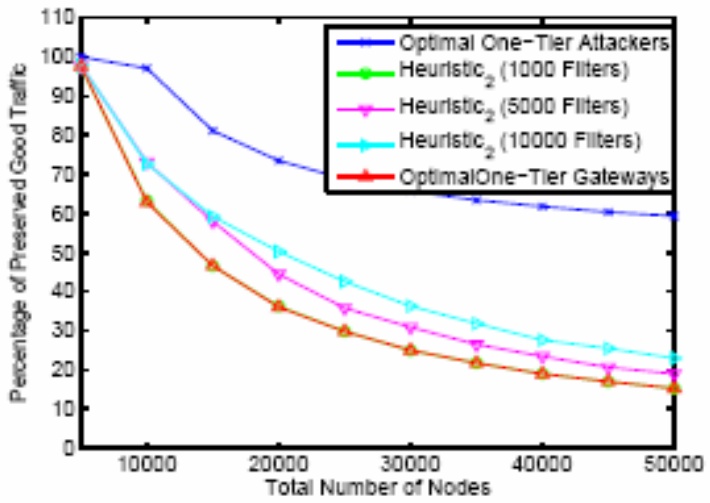

(b) Heuristic 2 (Gateway-based)

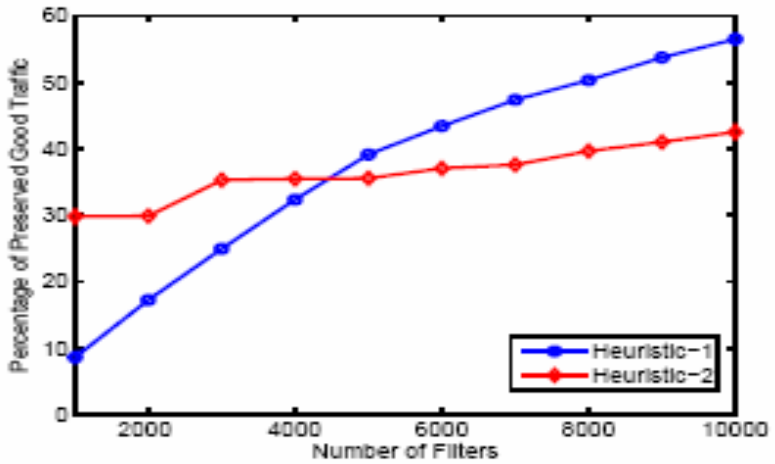

(c) Comparison of the two heuristics against number of filters

UCIrvine
University of California, Irvine

Attacker-based heuristic
for the uniform scenario

Heuristic vs. Optimal

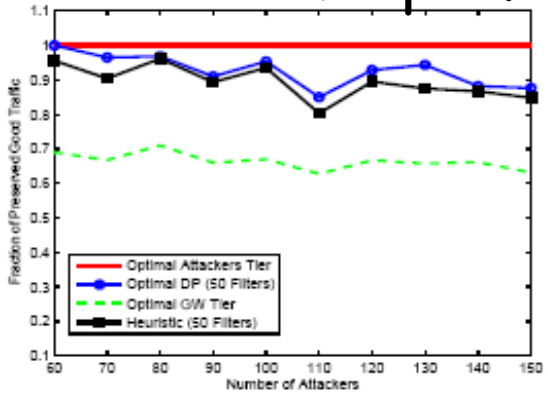

(a) Good traffic remains constant (60 sources), number of attackers increase (from 60 to 150 sources)

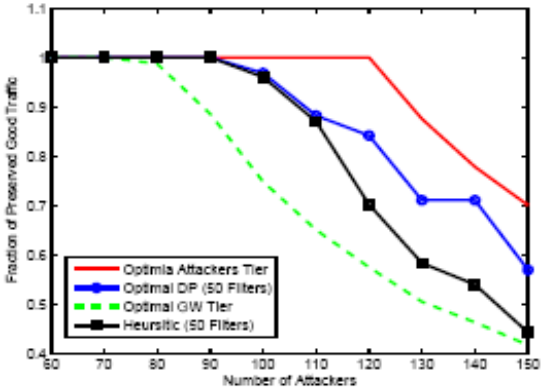

(b) Good traffic increases (from 10 to 100 sources) together with the attack traffic (from 60 to 150 sources)

Heuristic for more filters

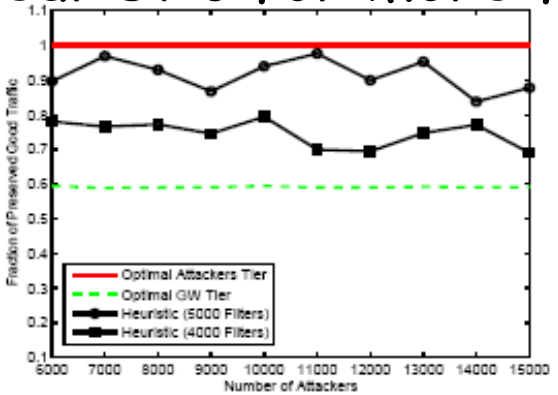

(a) Good traffic remains constant (6000 sources), number of attackers increase (from 6000 to 15000 sources)

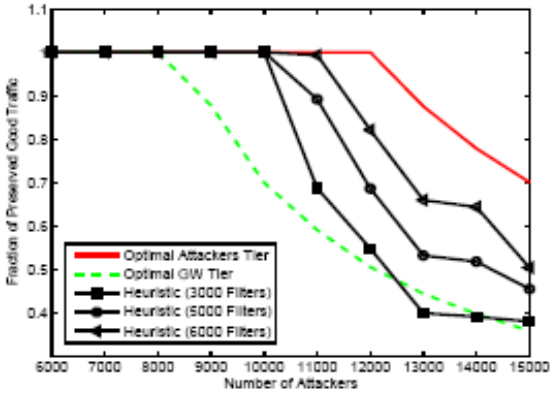

(b) Good traffic increases (from 1000 to 10000 sources) together with the attack traffic (from 6000 to 15000 sources)

Summary

- Formulated filtering as a resource allocation problem
- Showed that optimal filtering brings benefit in common attack scenarios
- Preliminary heuristics

- Impact
  - Optimal allocation provides a bound on filtering performance, under ideal assumptions
  - Heuristics can improve practical filtering policies

Ongoing Work

- Efficient filtering algorithms cont'd
- Evaluation with attack traces

- Imperfect identification of attackers
  - combine it with approaches for detection and/or reputation systems
- Filtering against Dynamic Attacks
  - complexity: incremental updates of the solution
  - adapting to adversarial attack distribution
- Filtering on a network
  - So far we considered only one router
  - coordination across several routers

Thank you!

athina@uci.edu

http://aegean.eng.uci.edu/