A Novel Framework for Modeling and Mitigating Distributed Link Flooding Attacks

Christos Liaskos[1] V. Kotronis[2] X. Dimitropoulos[1]

[1]Foundation of Research and Technology - Hellas (FORTH), Greece
[2]ETH Zurich, Switzerland
Emails: **cliaskos@ics.forth.gr**, vkotroni@tik.ee.ethz.ch, fontas@ics.forth.gr

Funding source: European Research Council, Grant Agreement no. 338402, project "NetVolution"

Scope & Motivation

- DDoS Link-Flooding attacks have great potential:
  - Deplete the bandwidth of certain network links, disconnecting entire domains—even countries—from the Internet.
- DDoS attacks are a reality:
  - Spamhaus (2013): 300 Gbit/s of malicious traffic upon the intended target [1].
- DDoS attacks are evolving in stealth:
  - Crossfire, Coremelt [2, 3]: Flood links indirectly, with seemingly legit traffic.

Scope: Define a framework to Model, Understand, and Expose evolved DDoS attacks.

FORTH-ICS

Key-requirements for Defense

M. Nikkhah, C. Dovrolis and R. Guerin, "*Why didn't my (great!) protocol get adopted?*", Proceedings of ACM HOTNETS, November 2015.

1. Deliver as promised:
   1. Expose Crossfire attacks. (Mitigation is context-specific).
2. Be thin & non-disruptive:
   1. Do not upset the network's operation.
3. Add value to existing network mechanisms!
   1. Don't start your own, independent path!

FORTH-ICS

Relevant Existing Network Mechanisms

Traffic Engineering (TE):

- Existing & critical network mechanism.
- Natural reaction to link-flooding events, regardless of cause.
- Accomplished in two phases:
  - Calculate optimal load per network path (*load-balancing*).
  - Map traffic flows to paths, upholding the optimal loads.
    - ***Note:*** *The mapping is done random*ly!
    - ***Key-idea:*** Optimize *flow mapping* for attack exposure.

FORTH-ICS

Our Contributions

- An analytical framework *to understand* Crossfire attacks.
  - Wide applicability (multigraphs, multipath routing, generic bot behavior).
- A thin and scalable way *to implement* the framework in practice.
  - All done in an SQL DB, standard SQL queries only.
  - SDN and NFV-compliant design.
- An open-source simulator *to experiment* with Crossfire attacks.

FORTH-ICS

Modeling the Crossfire Attack

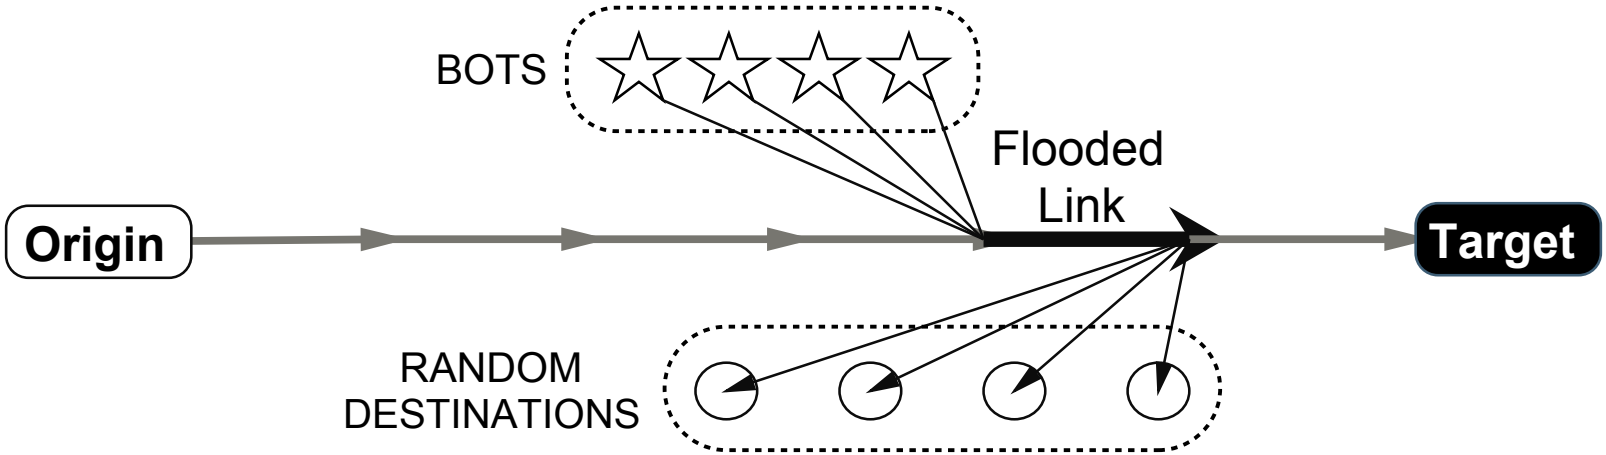

Modeling the Attack Cycle: Reactive Crossfire

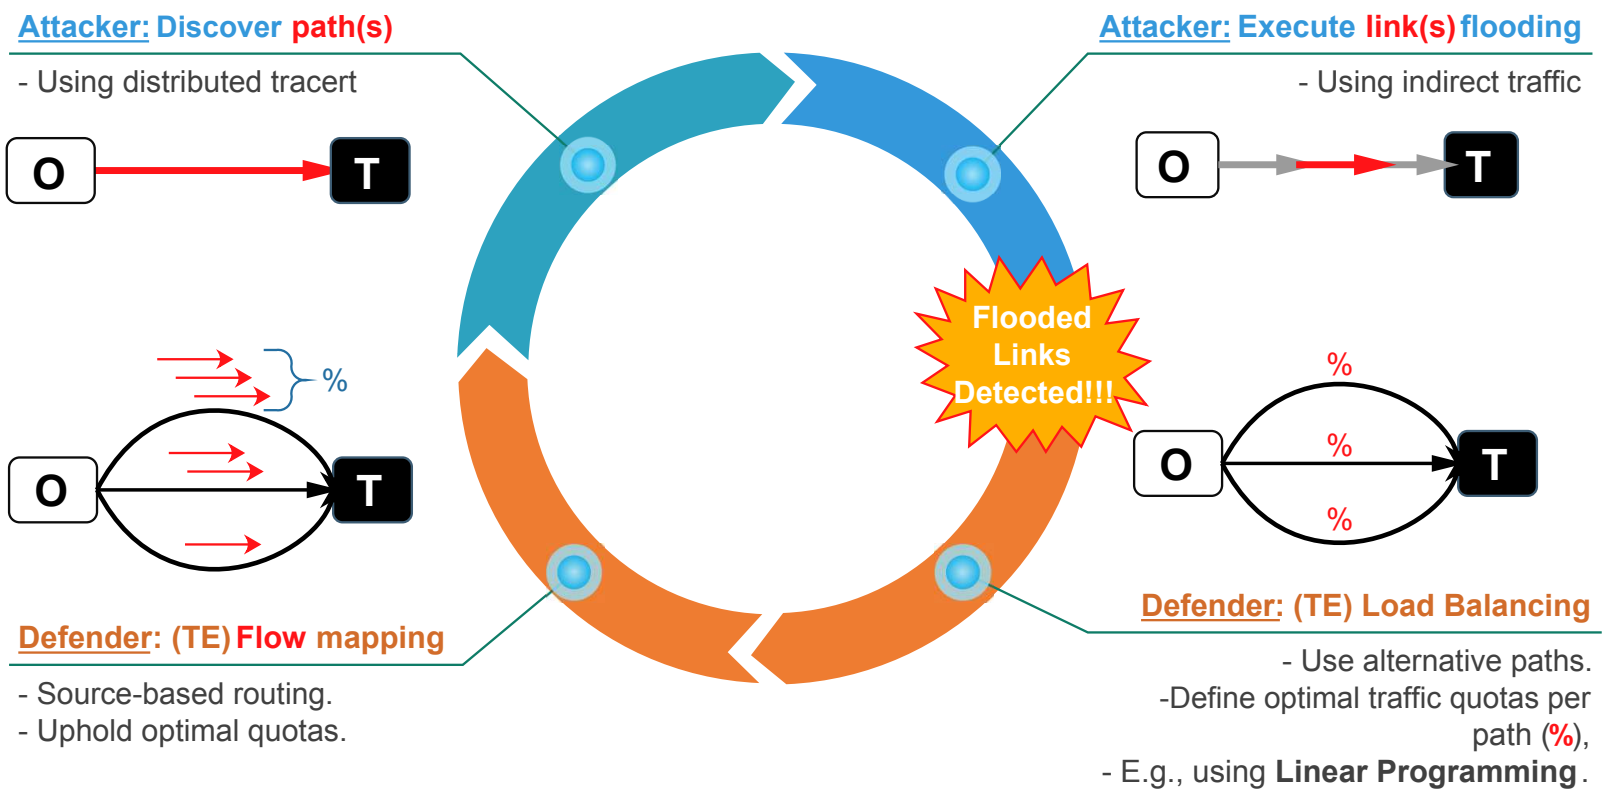

FORTH-ICS

Modeling the Attack Exposure with Associative Relations

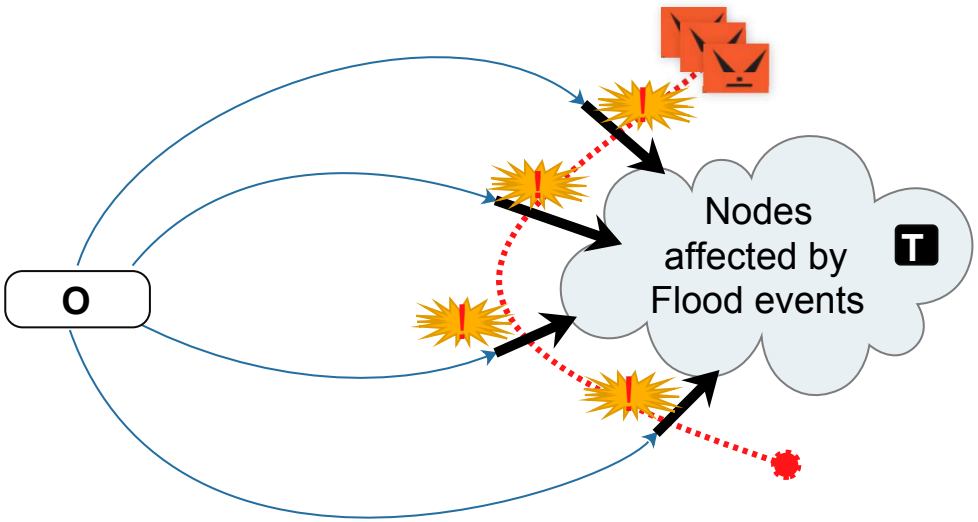

Relation at time t:

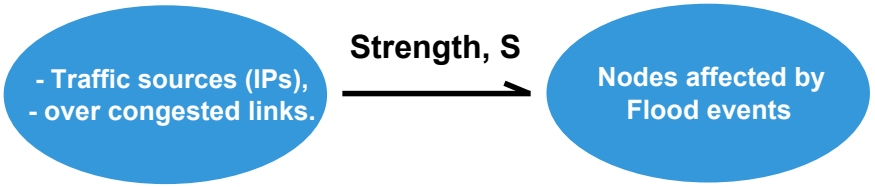

- An attack at cycle t floods a set of links, affecting nodes $\mathcal{N}(t)$.
  - $\mathcal{E}(t)$: Traffic sources (IPs) over congested links.

FORTH-ICS

Modeling the Attack Exposure with Associative Relations

- For all cycles 0...t, form the relation:
  $\mathscr{R}(t) : \underset{\forall t}{\cup}\mathscr{E}(t) \longrightarrow \underset{\forall t}{\cup}\mathscr{N}(t)$.
  - $e \xrightarrow{s} n$, *"entity e attacks node n"*.
  - $* \xrightarrow{s} n$, *"node n is an attack target"*.
  - $e \xrightarrow{s} *$, *"entity e is a bot"*.
- Strength s: # of observations for a relation (running).

FORTH-ICS

Effects of Attacker's Actions on the Detection Process

The attacker seeks to remain hidden.

- Hide the identity of bots and targets.

Ensures $\mathscr{E} \longrightarrow \mathscr{N}$ contains many false bots **e** and targets **n**.

Terminology:

- Increase $\|\mathscr{E}\|$, (Left-specificity).
- Increase $\|\mathscr{N}\|$, (Right-specificity).

FORTH-ICS

Effects of Load Balancing on the Detection Process

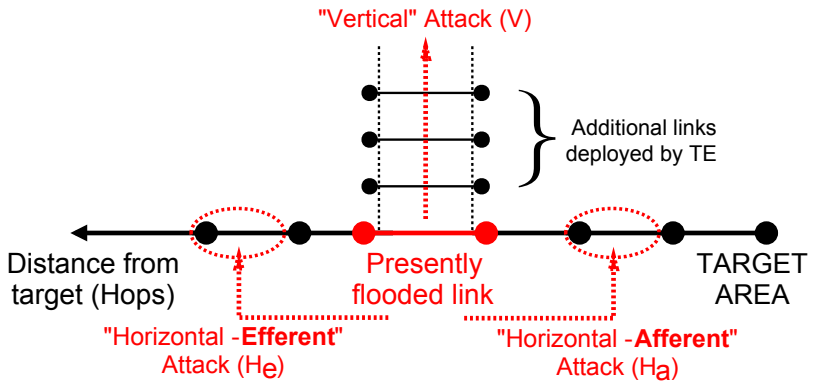

Figure: TE reclaims/adds capacity around congested areas. The attacker responds to keep affecting the target.

Table: Effects on the L / R specificity of observed relations.

Attacker's response types		
Vertical	Horizontal-efferent	Horizontal-afferent
$+, \pm$	$+, -$	$-, +$

FORTH-ICS

Optimizing TE Flow Mapping for Attack Exposure

Defender's goal

At t, reroute traffic flows such as, at t+1: $\min\{s\}\, e \xrightarrow{s} n, \forall r \subset \mathscr{R}(t)$

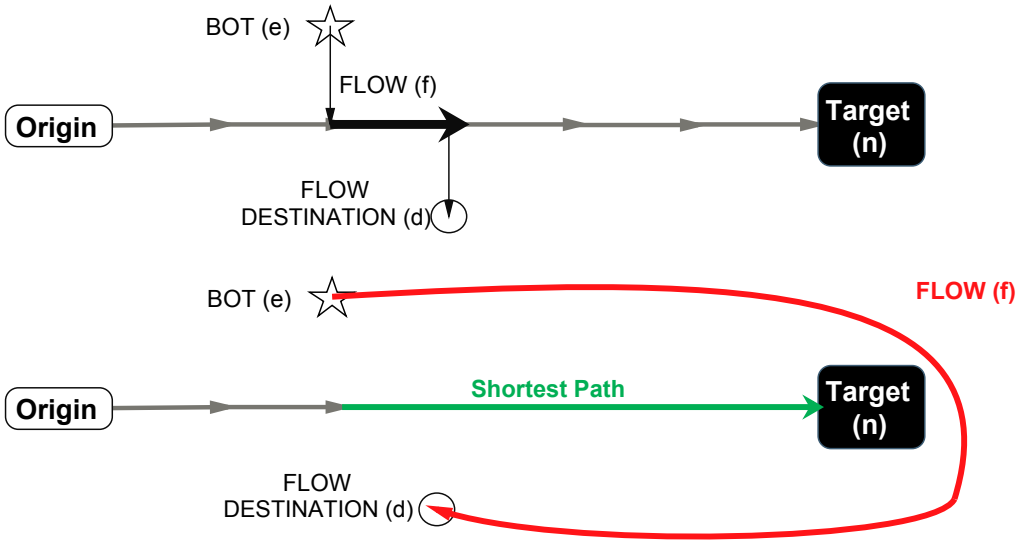

Qualitative meaning:

i) If flow retains destination $\xrightarrow{t+1}$ No attack, (RED is disjoint).

ii) If it changes $\xrightarrow{t+1}$ Minimal probability of accidental attack (GREEN).

FORTH-ICS

Practical Implementation of the Exposure Process

Intuitive implementation **via SQL**.

- Maintain two, simple tables, updated on link congestion events.

PROBABLE_BOTS

PK	Src_IP	Flooded link
1	31.32.4.1	L1
2	31.32.4.3	L1
3	31.32.4.4	L2
4	31.32.4.5	L2

PROBABLE_TARGETS

PK	NodeID	Flooded link
1	1	L1
2	2	L1
3	1	L2
4	3	L2

Relation exposure:

- **e**$\rightarrow$* : SELECT Src_IP, count(PK) AS strength from PROBABLE_BOTS GROUP BY Src_IP
- *$\rightarrow$**n** : SELECT NodeID, count(PK) AS strength from PROBABLE_TARGETS GROUP BY NodeID
- **e**$\rightarrow$**n** : ... INNER JOIN on FLOODED_LINK ...

FORTH-ICS

Simulations / Setup

We evaluate:

1. Effectiveness of attacker's bot/target *obfuscation attempts* (L/R specificity)
   1. Attack using some of the the bots only.
   2. Attack more than one targets.
   3. **Also:** Natural re-homing of legit flows.
2. The role of the *topology*.
   - 50 real ISP topologies (Internet Topology-Zoo).
   - *scenario:* cut-off an ISP POP from the Internet.

FORTH-ICS

Results I: Obfuscation Effects

Easier to obfuscate the target than the bots

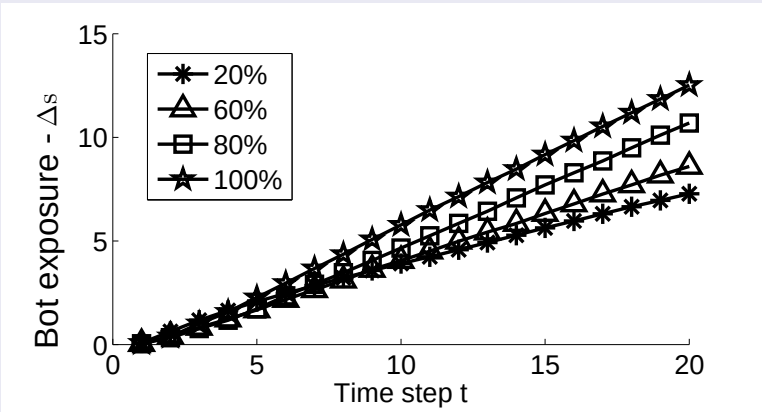

(a) L-specificity ($\varepsilon \rightarrow \star$ relations): Probabilistic bot participation to an attack. % $\rightarrow$allowed bot re-use.

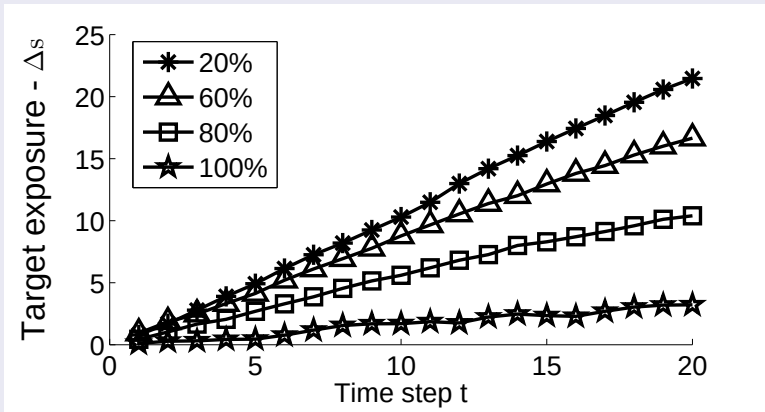

(b) R-specificity ($\star \rightarrow$ n relations): Effects of attacking random nodes. % is the flow rehome_ratio.

$\Delta s = E[s_{bots}] - E[s_{benign}]$
for $\epsilon \xrightarrow{s} \star$ relations

$\Delta s = E[s_{target\ node}] - E[s_{benign}]$
for $\star \xrightarrow{s} n$ relations

FORTH-ICS

Results II: Topological Effects

Decentralized topologies (i.e., not star-like): easier to attack, easier to detect.

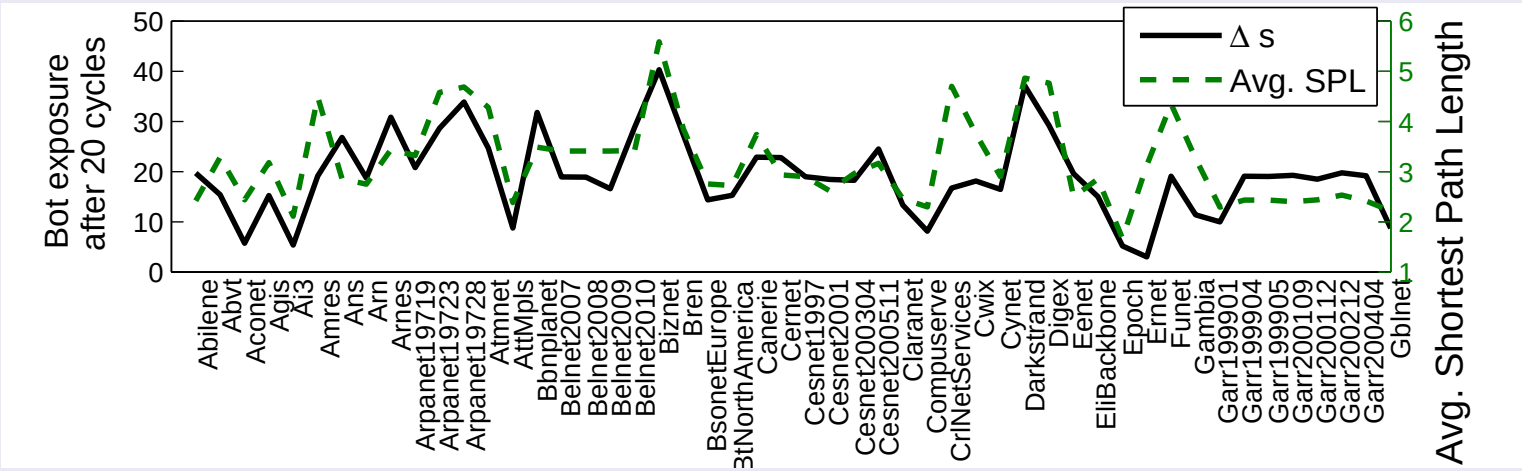

Figure: Effect of topology on the specificity of the $\varepsilon \rightarrow \star$ relations. Strong correlation to the average-shortest-path-length (avg. SPL) topology metric.

FORTH-ICS

Summary

- Introduced a novel framework for studying stealthy DDoS link-flooding attacks.
- **Goal:** Facilitate detection of susceptible bots and targets.
- Use relational algebra to formulate **bots** $\rightarrow$ **target** relations.
  - **Benefit:** Ease & scalability of implementation (SQL).
- **Entry point:** the TE process.
  - Use same inputs, leave TE load-balancing objective untouched.
  - Detection-optimal mapping of flows-to-paths.
  - **Key-idea:** keep probable bots targets @ separate paths, punish persistence.

FORTH-ICS

Outlook

- Build upon the analytical framework:
  - Express Attack/Defense *strategies* (Game-theory).
- Quantify the vulnerability of a given topology as a *metric*.
- Distribute the defense scheme as an SDN security app.
  - FRESCO framework.

Shin, Seungwon, et al.

FRESCO: Modular Composable Security Services for Software-Defined Networks. NDSS. 2013.

FORTH-ICS

Thank you!

- JAVA Simulator available at:
  - **http://users.ics.forth.gr/cliaskos/#PUBLICATIONS**

References


The DDoS That Almost Broke The Internet, (2014)
http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet.


Kang M. et al., *The Crossfire Attack,*
Proc. of Security & Privacy (SP'13).


Studer A. et al., *The Coremelt Attack,*
Proc. of ESORICS'09.

FORTH-ICS