Modeling and Mitigating the Coremelt Attack

Guosong Yang[1], Hossein Hosseini[2], Dinuka Sahabandu[2], Andrew Clark[3], João Hespanha[1], and Radha Poovendran[2]

[1]Department of Electrical and Computer Engineering,
University of California, Santa Barbara

[2]Department of Electrical Engineering,
University of Washington

[3]Department of Electrical and Computer Engineering,
Worcester Polytechnic Institute

2018 American Control Conference

Introduction

- The Coremelt attack on a TCP network with the “dumbbell” topology

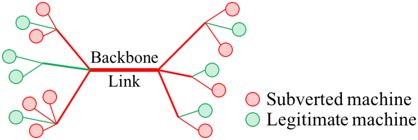

- Contribution
  - A dynamical system model for analysis
  - A limited number of subverted machines (bots): a modified TCP algorithm
  - A flow-based mitigation method
  - Simulation results

Distributed denial of service (DDoS) attack

- Attempt to disrupt network service by sending superfluous traffics from a vast number of bots

Distributed denial of service (DDoS) attack

- Attempt to disrupt network service by sending superfluous traffics from a vast number of bots
- Soaring number of Internet of Things (IoT) $\Longrightarrow$ Escalating DDoS threats
  - 21 billion IoT devices by 2020

Distributed denial of service (DDoS) attack

- Attempt to disrupt network service by sending superfluous traffics from a vast number of bots
- Soaring number of Internet of Things (IoT) $\Longrightarrow$ Escalating DDoS threats
  - 21 billion IoT devices by 2020
- One of world's largest DDoS attack to date [Ant+17]
  - 2016 on OVH (hosting service in France)
  - Mirai Botnet: 150,000 hacked IoT devices, 600,000 at peak
  - Attack flow rate: 1 Tbps

[Ant+17] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, in *26th USENIX Secur. Symp.*, 2017

The Coremelt attack

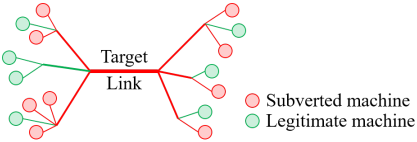

- A link-flooding DDoS attack [SP11]
- Target: backbone link

[SP11] A. Studer and A. Perrig, in *16th Eur. Symp. Res. Comput. Secur.*, 2011

The Coremelt attack

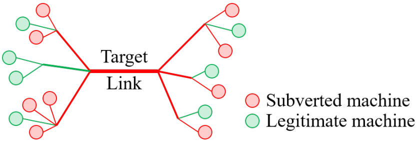

- A link-flooding DDoS attack [SP11]
- Target: backbone link
- Distributed botnet
  - Available
    - Mirai Botnet: 150k bots, 600k at peak
    - Among M bots there are $O(M^2)$ connections
  - Affordable
    - Price per 1000 bots: \$100–\$180 in U.S. or U.K., \$20–\$60 in Europe, less than \$10 elsewhere

[SP11] A. Studer and A. Perrig, in *16th Eur. Symp. Res. Comput. Secur.*, 2011

The Coremelt attack

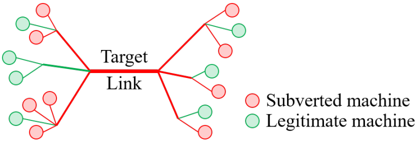

- A link-flooding DDoS attack [SP11]
- Target: backbone link
- Distributed botnet
  - Available
    - Mirai Botnet: 150k bots, 600k at peak
    - Among M bots there are $O(M^2)$ connections
  - Affordable
    - Price per 1000 bots: \$100–\$180 in U.S. or U.K., \$20–\$60 in Europe, less than \$10 elsewhere
- Low-intensity, legitimate-looking traffic
  - Able to evade conventional DDoS defenses

[SP11] A. Studer and A. Perrig, in *16th Eur. Symp. Res. Comput. Secur.*, 2011

Transmission Control Protocol (TCP)

- A congestion control algorithm [Pos81]
  - One congestion window per round-trip time (RTT)
  - Detect congestion based on missing acknowledgements (ACKs)
  - Additive-increase/multiplicative-decrease (AIMD) feedback algorithm [CJ89]

[Pos81] J. Postel, Information Sciences Institute, Tech. Rep., 1981
[CJ89] D.-M. Chiu and R. Jain, *Comput. Networks ISDN Syst.*, 1989

Transmission Control Protocol (TCP)

- A congestion control algorithm [Pos81]
  - One congestion window per round-trip time (RTT)
  - Detect congestion based on missing acknowledgements (ACKs)
  - Additive-increase/multiplicative-decrease (AIMD) feedback algorithm [CJ89]
- TCP-NewReno [Hen+12]
  - Widely used in modern Internet
  - Better for bursts of packet drops

[Pos81] J. Postel, Information Sciences Institute, Tech. Rep., 1981
[CJ89] D.-M. Chiu and R. Jain, *Comput. Networks ISDN Syst.*, 1989
[Hen+12] T. Henderson, S. Floyd, A. Gurtov, and Y. Nishida, Internet Engineering Task Force, Tech. Rep., 2012

Dynamical system model

- Analyze the impact and effectiveness of the Coremelt attack
- Establish flow composition and convergence via Lyapunov-based analysis
- Understand the relations between the number of bots, packet drop probability, and link usage ratio of users
- Develop a flow-based mitigation method

Network model

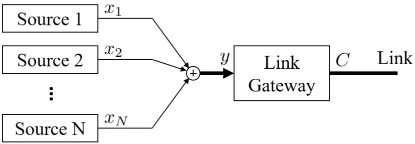

TCP-NewReno source

- One congestion window w_k per RTT τ_k
- Average flow rate $x_k = w_k/\tau_k$
- Congestion probability $q_k \approx w_k p$ with packet drop probability p

Network model

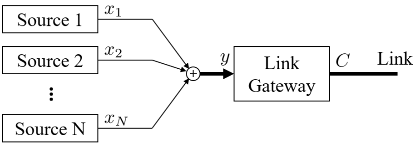

TCP-NewReno source

- One congestion window w_k per RTT τ_k
- Average flow rate $x_k = w_k/\tau_k$
- Congestion probability $q_k \approx w_k p$ with packet drop probability p
- AIMD algorithm for TCP-NewReno

$$\begin{cases} w_k \leftarrow w_k + 1, & \text{without congestion;} \\ w_k \leftarrow w_k/2, & \text{with congestion} \end{cases}$$

Network model

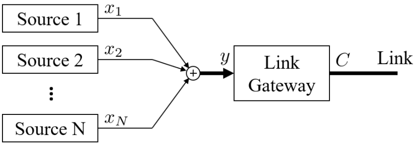

TCP-NewReno source

- One congestion window w_k per RTT τ_k
- Average flow rate $x_k = w_k/\tau_k$
- Congestion probability $q_k \approx w_k p$ with packet drop probability p
- AIMD algorithm for TCP-NewReno

$$\begin{cases} w_k \leftarrow w_k + 1, & \text{without congestion;} \\ w_k \leftarrow w_k/2, & \text{with congestion} \end{cases}$$

- Dynamical system model:

$$\dot{x}_k = \frac{1}{\tau_k^2}\left((1 - q_k) - \frac{w_k}{2} q_k\right)$$

Network model

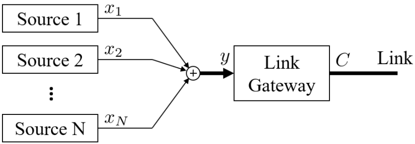

TCP-NewReno source

$$\dot{x}_k = \frac{1 - \tau_k x_k p}{\tau_k^2} - \frac{p x_k^2}{2}, \qquad k = 1, \ldots, N$$

Network model

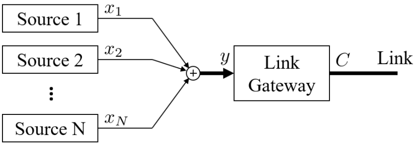

TCP-NewReno source

$$\dot{x}_k = \frac{1 - \tau_k x_k p}{\tau_k^2} - \frac{p x_k^2}{2}, \qquad k = 1, \ldots, N$$

Bottleneck link

- Aggregate rate $y = \sum x_k$
- Bandwidth C
- Drop the excess packets

$$p = \begin{cases} 1 - C/y, & \text{if } y > C; \\ 0, & \text{otherwise} \end{cases}$$

Attack with M bots following TCP-NewReno

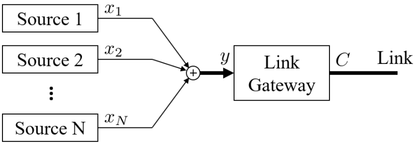

Attack with M bots following TCP-NewReno

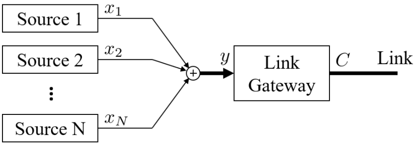

Theorem 1

- If M bots and $N - M$ users all follow TCP-NewReno, the dynamical system is globally asymptotically stable (GAS)
- Packet drop probability converge to p^* satisfying $\sum_{k=1}^{N} \frac{1}{\tau_k} = \frac{\sqrt{1+2/p^*}+1}{2(1-p^*)} p^* C$

Attack with M bots following TCP-NewReno

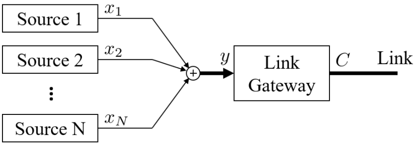

Theorem 1

- If M bots and $N-M$ users all follow TCP-NewReno, the dynamical system is globally asymptotically stable (GAS)
- Packet drop probability converge to p^* satisfying $\sum_{k=1}^{N} \frac{1}{\tau_k} = \frac{\sqrt{1+2/p^*}+1}{2(1-p^*)} p^* C$

Proof

- Lyapunov function $V(x-x^*)$ such that
$$\dot{V}(x-x^*) \leq -W(x-x^*) - (p-p^*)(y-y^*)$$
- $W(x-x^*)$ is positive definite
- Packet drop probability p is increasing in aggregate rate y

Attack with M bots following TCP-NewReno

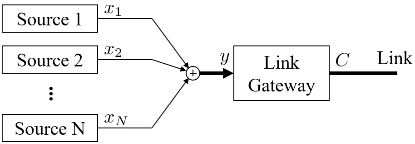

Theorem 1

- If M bots and $N-M$ users all follow TCP-NewReno, the dynamical system is globally asymptotically stable (GAS)
- Packet drop probability converge to p^* satisfying $\sum_{k=1}^{N} \frac{1}{\tau_k} = \frac{\sqrt{1+2/p^*}+1}{2(1-p^*)} p^* C$

Implication

- For the same RTT τ, the link usage ratio of users is $1 - M/N$

Attack with M bots following TCP-NewReno

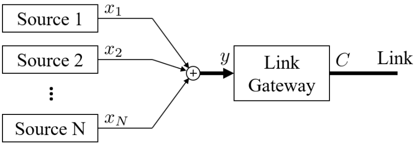

Theorem 1

- If M bots and $N-M$ users all follow TCP-NewReno, the dynamical system is globally asymptotically stable (GAS)
- Packet drop probability converge to p^* satisfying $\sum_{k=1}^{N}\frac{1}{\tau_k}=\frac{\sqrt{1+2/p^*}+1}{2(1-p^*)}p^*C$

Implication

- For the same RTT τ, the link usage ratio of users is $1-M/N$
- A target value p^* can be achieved by enough bots so that $N\geq\frac{\sqrt{1+2/p^*}+1}{2(1-p^*)}p^*\tau C$

Attack with M bots following a modified TCP

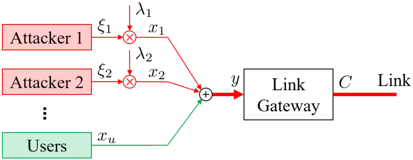

Attack with M bots following a modified TCP

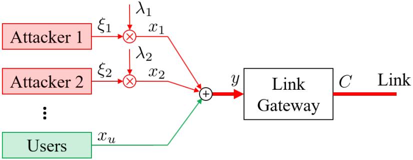

Modified TCP source

- Internal state ξ_j that follows the AIMD algorithm for TCP-NewReno
- Flow rate $x_j = \lambda_j \xi_j$ with gain $\lambda_j \geq 0$
- Drive the congestion probability to target value q_0 by slowly adjusting λ_j:

$$\dot{\lambda}_j = \gamma_j \xi_j (q_0 - q_j)^+_{\lambda_j}$$

Attack with M bots following a modified TCP

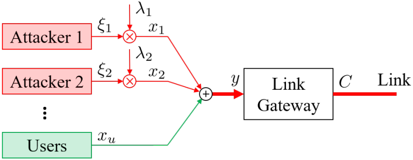

Theorem 2

- If $N - M$ users follow TCP-NewReno and M bots follow the modified TCP, the dynamical system is GAS
- Congestion probability converge to target value q_0 for any M

Attack with M bots following a modified TCP

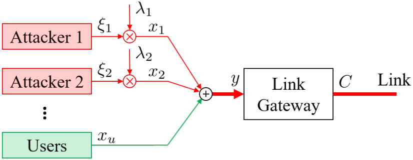

Theorem 2

- If $N - M$ users follow TCP-NewReno and M bots follow the modified TCP, the dynamical system is GAS
- Congestion probability converge to target value q_0 for any M

Proof

- Weak Lyapunov function $V(x_u - x_u^*, \xi - \xi^*, \lambda - \lambda^*)$ such that
$$\dot{V}(x_u - x_u^*, \xi - \xi^*, \lambda - \lambda^*) \leq -W(x_u - x_u^*, \xi - \xi^*) - (p - p^*)(y - y^*)$$
- $W(x_u - x_u^*, \xi - \xi^*)$ is positive definite, p is increasing in y
- LaSalle's invariance principle

Mitigation

- Detection-based mitigation: source authentication, packets inspection
  - Less effective against Coremelt:
    - Communication between bot pairs
    - Low-intensity, legitimate-looking traffic

Mitigation

- Detection-based mitigation: source authentication, packets inspection
  - Less effective against Coremelt:
    - Communication between bot pairs
    - Low-intensity, legitimate-looking traffic
- Flow-based mitigation: penalize aggressive sources
  - Monitor source flow rates and assign individual drop probability p_k so that the bandwidth C is evenly shared: $p_k \sim 1 - C/(Nx_k)$

Mitigation

- Detection-based mitigation: source authentication, packets inspection
  - Less effective against Coremelt:
    - Communication between bot pairs
    - Low-intensity, legitimate-looking traffic
- Flow-based mitigation: penalize aggressive sources
  - Monitor source flow rates and assign individual drop probability p_k so that the bandwidth C is evenly shared: $p_k \sim 1 - C/(Nx_k)$
  - Advantages:
    - Guaranteed link usage ratio of users: $1 - M/N$
    - Does not require modifying source transmission protocols

Mitigation

- Detection-based mitigation: source authentication, packets inspection
  - Less effective against Coremelt:
    - Communication between bot pairs
    - Low-intensity, legitimate-looking traffic
- Flow-based mitigation: penalize aggressive sources
  - Monitor source flow rates and assign individual drop probability p_k so that the bandwidth C is evenly shared: $p_k \sim 1 - C/(Nx_k)$
  - Advantages:
    - Guaranteed link usage ratio of users: $1 - M/N$
    - Does not require modifying source transmission protocols
  - Limitations:
    - Extra resources needed to monitor source flow rates
    - Users with smaller RTTs will also be penalized
    - No effect against attacks with bots following TCP-NewReno

Simulation: without mitigation

- Network of $2,000$ users and $1,000$ bots
- Link capacity of 1 million packets per RTT

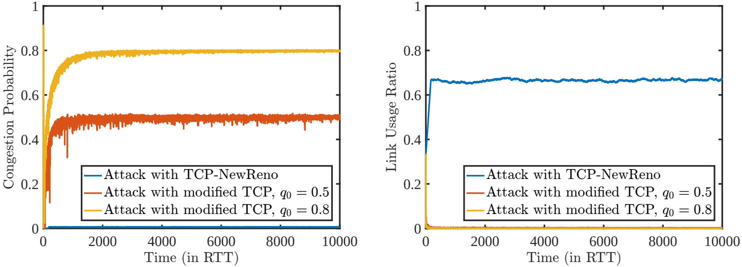

- Attack with TCP-NewReno: low congestion probability; link usage ratio of users is $2/3$
- Attack with modified TCP: target congestion probability; link usage ratio of users is low

Simulation: with mitigation

- Network of 2000 users and 1000 bots
- Link capacity of 10^6 packets per RTT

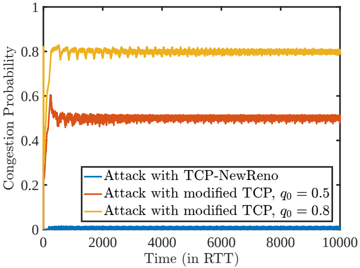

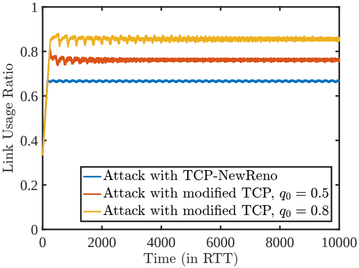

- Attack with modified TCP: target congestion probability; link usage ratio of users is high

Conclusion

- Contribution
  - A dynamical system model for analyzing the Coremelt attack on a TCP network
  - A limited number of bots: a modified TCP algorithm
  - A flow-based mitigation method
  - Simulation results

Conclusion

- Contribution
  - A dynamical system model for analyzing the Coremelt attack on a TCP network
  - A limited number of bots: a modified TCP algorithm
  - A flow-based mitigation method
  - Simulation results
- Future work
  - User Datagram Protocol (UDP) [Pos80]
  - The Crossfire attack [KLG13]

[Pos80] J. Postel, Information Sciences Institute, Tech. Rep., 1980
[KLG13] M. S. Kang, S. B. Lee, and V. D. Gligor, in *2013 IEEE Symp. Secur. Priv.*, 2013

References

[Ant+17] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, "Understanding the Mirai botnet," in *26th USENIX Secur. Symp.*, 2017.

[CJ89] D.-M. Chiu and R. Jain, "Analysis of the increase and decrease algorithms for congestion avoidance in computer networks," *Comput. Networks ISDN Syst.*, 1989.

[Hen+12] T. Henderson, S. Floyd, A. Gurtov, and Y. Nishida, "The NewReno Modification to TCP's Fast Recovery Algorithm," Internet Engineering Task Force, Tech. Rep., 2012.

[KLG13] M. S. Kang, S. B. Lee, and V. D. Gligor, "The Crossfire attack," in *2013 IEEE Symp. Secur. Priv.*, 2013.

[Pos80] J. Postel, "User Datagram Protocol," Information Sciences Institute, Tech. Rep., 1980.

[Pos81] J. Postel, "Transmission Control Protocol," Information Sciences Institute, Tech. Rep., 1981.

[SP11] A. Studer and A. Perrig, "The Coremelt attack," in *16th Eur. Symp. Res. Comput. Secur.*, 2011.

Acknowledgements