On the Potential Threats of "Smart" Traffic Flooding Attacks

(Poster version: http://www.seas.upenn.edu/~yingx/Sigcomm2004Poster.pdf)

Ying Xu and Roch Guérin *
Multimedia and Networking Lab.
Department of Electrical and Systems Engineering
University of Pennsylvania
Philadelphia, PA, USA
yingx,guerin@ee.upenn.edu

Network Denial-of-Service (DoS) attacks can severely disrupt regular service delivery in the Internet. The substantial threats posed by DoS attacks have triggered the development of many *router-based* defense solutions. For example, *reactive* DoS defense systems such as the ACC-Pushback [1] proposal detect an ongoing attack based on monitoring *link loss rate*. If the measured loss rate exceeds a certain threshold, the ACC system is activated and starts searching for malicious traffic. *Proactive* defense systems have also been designed that continuously monitor the behavior of all users for conformance with a given criterion. For instance, the RED-PD proposal [2] tries to enforce TCP friendliness by sampling the incoming traffic using random packet losses. Any flow that loses more than the expected number of losses of a standard TCP flow during a certain period is deemed non-conformant and will be regulated via packet filtering. Since typical DoS attacks nowadays are mostly high rate attacks that flood the target host or network domain, they are usually accompanied with heavy losses and are likely to be identified and contained by either type of defense.

We expect that, as with other kind of "attacks" such as worms and viruses, DoS attacks will evolve to overcome defenses. Understanding if and when existing defenses can be defeated by more intelligent attackers, and what the implications are in designing better defense mechanisms is, therefore, an important topic. To gain a basic insight into this issue, we developed a number of "smarter" attacking schemes along what we felt were the most natural and promising directions, and evaluated their impact on existing defense schemes such as ACC and RED-PD. We explored two dimensions along which the intelligence (and complexity) of an attacking scheme increases. The first dimension is related to the level of sophistication of a *single* attacker. Specifically, we considered attackers that target either bandwidth or buffer space as the resource they are attempting to deprive other users from without being detected. When a single attacker was not successful, we turned to scenarios involving multiple attackers or more precisely, attackers with *multiple identities*. This accounts for cases when attackers have recruited multiple daemon hosts, or when a single attacker disguises itself using multiple source addresses. The overall complexity of an attacking scheme is then a function of the the number of entities involved, the complexity of each attacking entity, and the level of coordination between them. Our goal is to investigate whether the combination of increased intelligence and the use of multiple identities, can translate into greater attacking efficiency at a reasonable cost. Our main results are as follows:

*This work was supported in part through NSF grant ITR00-85930.

SIGCOMM 2004, Aug.30–Sep.3, 2004, Portland, OR, USA.

Single Attacker Scheme: We found that a *single* attacker employing a *rate adaptation* strategy can successfully evade the ACC system. This type of attacker *slowly* increases its rate and thus the system congestion level, until the link loss rate *reaches and stabilizes* at a level below the ACC activation threshold. This allows the attacker to greatly degrade the performance of TCP users without triggering the underlying defense. However, because of its high intrinsic bandwidth consumption and correspondingly large number of lost packets, a single such attacker will be detected by the RED-PD system. Hence, we considered another type of attacker that instead tries to increase the losses of TCP users while minimizing its own, by attempting to track the evolution of the target link queue, and alternating between bursts and silences in order to fill the queue and then stop before experiencing significant losses. We found that for a RED queue, such a periodic "blasting" strategy can *sometimes* succeed in making TCP users experience loss rates an order of magnitude larger than the attacker. However, the *absolute* amount of packet losses of the attacker is typically still high enough for the RED-PD system to detect and throttle it.

Multiple Identities Scheme: Since it is hard for a single attacker to elude the RED-PD defense, we further examined attacking schemes involving multiple attackers. We found that the attacking efficiency in this case is greatly influenced by how the multiple attacking entities were selected. When attackers are traffic sources residing on different host machines, the number of attackers required to launch a successful attack is large, primarily because it is hard to achieve efficient coordination among sources. In contrast, when the multiple attackers instead corresponded to multiple identities, i.e., multiple source addresses, of a single host, and were selected *in turn* during an attack, a much smaller number of attackers (distinct addresses) was needed to foil the RED-PD system. This is due to the sampling mechanism used by the RED-PD system, which relies on samples within only one scanning interval to detect a TCP-unfriendly user.

Our investigation has revealed that by adapting their behavior, DoS attackers can indeed defeat existing defense systems. We believe that this understanding is a first step in allowing us to devise more efficient and robust DoS defense mechanisms, which is the topic of our current research.

1. REFERENCES

[1] R. Mahajan, S. M. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker. Controlling high bandwidth aggregates in the network. *SIGCOMM Comput. Commun. Rev.*, 32(3):62–73, 2002.

[2] R. Mahajan and S. Floyd. Controlling high-bandwidth flows at the congested router. In *Proc. 9th Intl. Conf. Netw. Protocols (ICNP'01)*, pages 192–201. IEEE Computer Society, 2001.