The Dynamics and Control of Internet Attacks

ELIZABETH BRADLEY AND JAMES G. GARNETT

Department of Computer Science
University of Colorado
Boulder CO 80309-0430 USA
lizb@cs.colorado.edu

The design of the internet did not anticipate malicious behavior. This lack of security in the architecture has rendered this vital system vulnerable to misuse by so-called 'malware:' viruses, spam, spyware, etc. A particularly unpleasant example of this is the denial-of-service (DoS) attack, in which an attacker bombards a victim—a webserver or router, for instance—with requests for service. In a more-pernicious variant called the distributed DoS (DDoS) attack, an attacker uses a virus to co-opt "zombie" machines, installing code on each that bonds them into a "botnet" that can generate a distributed, coordinated bombardment. In either case, when the barrage clogs the victim's input buffers, it cannot serve other customers and business is disrupted. Attacks like this began doubling every year or two in the late 80s, and this growth continues to accelerate. All of the major webservers, like amazon.com, have been affected, and an associated extortion racket has even emerged in the past few years; in a CMU study, 17% of businesses surveyed had been threatened with such an attack unless they paid protection money [1]. The costs of these attacks—outright loss, post-attack repair, and defense—have been estimated in the billions of dollars.

A variety of strategies have been developed to defend against DoS attacks, but attackers find ways to work around them within hours or days of their deployment. A strategy that identifies bursts of requests, for instance, can be fooled by reducing the attack flow to a trickle that flies just under the burst-detector's radar screen. This kind of co-evolution vastly complicates the task of recognizing and defending against DoS attacks. Note, too, that any defense strategy that examines individual requests in any way—e.g., identifying the IP address that is the source of the barrage—is doomed to fail; the attacker need only escalate the attack to exceed the rate at which requests can be examined in order to saturate the defender.

Viewing the internet as a complex, distributed, nonlinear adaptive system is an effective way to understand the dynamics of DoS attacks and develop viable defenses against them. In this talk, we discuss how to build stochastic models of the behavior of a resource like a webserver or router that is under attack. We then show how one can use such a model as the core of a nonlinear adaptive model-reference controller that allows that resource survive an attack gracefully, and with mathematically guaranteed performance. In lab bench tests, this controller allowed a server that was under DoS attack on one port to successfully service 100% of the valid requests on another port. (Without the controller, 97% of those requests were dropped.) It is the topic of a US patent, and is currently deployed in commercial webserver hardware.

References

[1] D. Talbot, "The Internet is Broken," *Technology Review*, December 2005.

Keywords: internet, dynamics, control