ISSN 2394-3777 (Print)
ISSN 2394-3785 (Online)
Available online at www.ijartet.com

International Journal of Advanced Research Trends in Engineering and Technology (IJARTET)
Vol. 2, Issue 2, February 2015

LRDOS Attacks in Internet Based Services

Zinniah.F[1], Hanitha.A[2], Andrew Rayan.T[3]
Department Of Information Technology, Francis Xavier Engineering College, Tirunelveli, India[1]
Department Of Information Technology, Francis Xavier Engineering College, Tirunelveli, India[2]
Assistant Professor, Department Of Information Technology, Francis Xavier Engineering College, Tirunelveli, India[3]

Abstract— Feedback control is a critical element in many Internet services (e.g., quality-of-service aware applications). Recent research has demonstrated the vulnerability of some feedback-control based applications to low-rate denial-of-service (LRDoS) attacks, which send high-intensity requests in an ON / OFF pattern to degrade the victim's performance and evade the detection designed for traditional DoS attacks. However, the intricate interaction between LRDoS attacks and the feedback control mechanism remains largely unknown. In this paper, we address two fundamental questions: 1) what is the impact of an LRDoS attack on a general feedback-control based system and 2) how to conduct a systematic evaluation of the impact of an LRDoS attack on specific feedback-control based systems. To tackle these problems, we model the system under attack as a switched system and then examine its properties. We conduct the first theoretical investigation on the impact of the LRDoS attack on a general feedback control system. We formally show that the attack can make the system's steady-state error oscillate along with the attack period, and prove the existence of LRDoS attacks that can force the system to be far off the desired state. In addition, we propose a novel methodology to systematically characterize the impact of an LRDoS attack on specific systems, and apply it to a web server and an IBM Notes server. This investigation obtains many new insights, such as new attack scenarios, the bound of the system's states, the relationship between the bound and the LRDoS attacks, the close-formed equations for quantifying the impact, and so on. The extensive experimental results are congruent with the theoretical analysis

Index Terms— Feedback control, low-rate DoS attack, switched system stability, performance degradation.

I. Introduction

FEEDBACK control is a fundamental building block of many Internet services. A classic example is the performance controller in a web server, which adjusts the server's configuration (e.g., admission rate) in response to the difference between the current and desired states for meeting expected performance (e.g., throughput and service time). Feedback control is also a central element in QoS-aware systems (e.g., cloud computing, high-performance computing , virtualized servers , cyber-physical systems , and autonomic computing .Recent studies have demonstrated that Low-Rate DoS (LRDoS) attacks can degrade the performance of some feedback-control based applications (instead of continuous) high-volume requests to force the victim away from the desired state, thus deteriorating its performance. Moreover, LRDoS attacks can escape the detection designed for flooding-based DoS attacks because of their ON/OFF traffic patterns. A burst of requests can be termed an attack pulse, and an LRDoS attack then consists of a sequence of attack pulses. Although seminal studies have showed the possibility of using LRDoS to attack feedback-control based systems, and have studied the corresponding damage to a few applications under limited attack patterns , they have not solved two fundamental questions. (1) What is the impact of an LRDoS attack on a general feedback control system? (2) How to conduct a systematic evaluation of the impact of an LRDoS attack on specific feedback-control based systems? It is challenging to tackle these questions because LRDoS attacks can force a victim system to exhibit discontinuous behavior on the arrival of attack pulses. Traditional control theory cannot handle such situation because it targets a pure continuous (or discrete) system represented by differential (or difference) equations . In this paper, to address the questions, we propose to model the system under attack as a switched system, which is a hybrid system composed of several subsystems and a switching law that indicates the sequence of subsystems .We first model the impact of an LRDoS attack on a general feedback control system in two important respects (Section III): steady-state error and system state .We prove the existence of LRDoS attacks that can constrain the victim system to a state, which is determined by an attacker and diverges away from the desired state, by proving that the system under attack can still be Lyapunov and Lagrange stable. The investigation of the general feedback control system motivates us to develop a novel methodology to systematically analyze the impact of an LRDoS attack on specific feedback-control based systems. We apply our methodology to two specific systems: the web server and the feedback-control based IBM Notes server .

II. Related Work

DDoS attacks have been plaguing the Internet for decades. They drain bandwidth and/or system resources to prevent normal users from receiving quality service . Traditional flood-based attacks can be easily detected because of their continuously high sending rates . In contrast, low-rate DoS attacks polymorphic traffic patterns and low average sending rates .

A.Low-Rate DoS Attacks

LRDoS attacks was first proposed to throttle the throughput of TCP connections by causing intermittent packet losses . Zhang et al. and Schuchard et al. showed that an attacker can launch LRDoS attacks on BGP sessions for crippling the Internets control plane. Recently, researchers examined the vulnerability of other applications to LRDoS attacks, including Internet services , load balancers , wireless networks , and peer-to-peer networks .Guirguis et al. found that an LRDoS attack can force a feedback a web server . There are three major differences between our work and theirs. First, while Guirguis et al. described the possibility of launching the LRDoS attack on a feedback-control based system, we formally show that the LRDoS attack can compel a feedback control system to stay at a state other than the desired state by proving that the system under attack is Lyapunov and Lagrange stable. Second, we propose a novel methodology to systematically evaluate the impact of an LRDoS attack on specific systems. This methodology enables us to obtain many new insights that are not reported before. For example, for the same web server, we reveal in Section IV that an attacker can launch three types of LRDoS attacks by adjusting the attack period. Guirguis et al. only examined the third type of LRDoS attack [16]. It is worth noting that the other two types of attacks can cause severer damage to the web server, as shown in Fig. 4. Moreover, we thoroughly analyze each type of LDRoS attack, including giving closed-form expressions for the throttled admission rate, determining the conditions under which the LRDoS attack will make the web server Lyapunov and Lagrange stable, deciding the bound of the system's state and the relationship between the bound and the LRDoS attack, and identifying the relationship between the damage caused by an LRDoS attack and its cost. Third, we employ our methodology to analyze the vulnerability of a feedback-control based IBM Notes server to the LRDoS attack. Note that the IBM Notes server is different from the web server in three aspects, including the application (i.e., email service vs. web service), the feedback controller (i.e., I controller vs. PI controller), and the system model.

Our analysis formally shows that an LRDoS attack can cause severe damage to the IBM Notes server. The investigation of these two kinds of servers provides convincing evidences on the threat of LRDoS attacks to feedback-control based systems and the generality of our methodology. Maciá-Fernández et al. proposed a smart attack called Low Rate DoS attack against Application Servers (LoRDAS), which dispatches attack requests to the victim server at carefully selected instances to occupy the server's queue and consequently prevent it from serving legitimate requests. They also built a mathematical model for LoRDAS attacks and evaluated their performance. The major difference between our work and theirs is that the attack examined in our paper exploits the feedback control mechanism in the victims system. However, LoRDAS attack takes advantage of the victim's queue. Moreover, an LoRDAS attack needs to predict the time instants when the queue of the victim server has free position so that it can make the attack requests reach the server around these time instants , whereas an LRDoS attack can send attack pulses with fixed or randomized interval to the victim system.

B. Defending Against LRDoS

LRDoS attacks have ON/OFF traffic patterns, they can evade detection schemes targeting flooding-based DoS attacks and therefore have motivated the design of new detection approaches. However, these approaches cannot be directly used to detect LRDoS attacks against Internet services for two reasons. First, as all of these approaches aim at LRDoS attacks targeting TCP or other systems (e.g., wireless networks, P2P networks, etc.), they rely on features specific to TCP and those systems. For example, we proposed the detection of anomalies in incoming TCP data traffic and outgoing TCP ACK traffic . Shevtekar et al. regarded a TCP flow as malicious if its period is equal to the fixed minimal RTO and its burst length is no less than other connections' RTTs . To detect LRDoS attacks using spoofed IP addresses, Shevtekar et al. captured anomalies that short-lived flows occupy a high percentage of the total traffic going through a link [40]. We proposed a new metric named the congestion participation rate (CPR) to infer attack flows that try to send more packets during congestion . To detect distributed LRDoS attacks, Xiang et al. used generalized entropy and information distance to quantify anomalies in packets, and required the control of all routers in the network. However, the detection of LRDoS attacks aimed at Internet services requires new metrics .

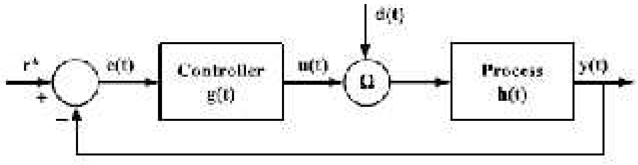

Fig. 1. A general feedback control system

Second, the majority of the previous work focuses on the Shrew attack that has a fixed attack period equal to TCP minimal RTO. For example, the spectral-analysis approach relies on the traffic spectrum of Shrew attack flows, which is different from that of normal flows because of the even attack period. However, LRDoS attacks can change their attack periods for mimicking normal flows. Sun et al. suggested using autocorrelation and dynamic time warping (DTW) to detect Shrew attacks, because their traffic bursts are the same and have fixed periods . However, it is unnecessary for LRDoS attacks to have invariable periods and similar attack pulses .

III. Attacking A General Feedback Control System

In this section, we investigate the impact of an LRDoS attack on a general feedback control system as shown in Fig. 1.This system comprises two major components: a process (i.e.,h(t)) and a controller (i.e., g(t)). h(t) represents any Internet service (e.g., web service, video streaming, etc.) while g(t) generates a control signal (i.e., u(t)) to regulate h(t) [2]. The input to the controller is a control error (i.e., e(t)), which is the difference between h(t)'s output (i.e., y(t)) and the expected value r ∗ . y(t) can be any measurable metric, such as system utilization or queue length. r ∗ is usually selected for the system to achieve the best performance by the system designer, and the controller drives y(t) towards r ∗ based on e(t). d(t) denotes the arrival rate of requests. To simplify the discussion, we assume that the arrival rate of normal requests is constant, denoted as λn . We compare the results when d(t) is constant or random for two real systems in Section VI. Depending on the application, d(t) enters into the process through an additive or a multiplicative operator . The steady-state error, defined as e s (t) = lim t →∞ (e(t)), is the first performance index of a feedback control system [18]. As e s (t) quantifies the accuracy of control, it should be as small as possible and preferably zero. Since all practical feedback control systems are necessarily stable , we assume that the victim system is Lyapunov stable without attack, meaning that the system trajectory can be kept arbitrarily close to an equilibrium point by starting sufficiently close to it . We also assume that the control error, the output, and the control signal induced by an attack pulse are not equal to zero. Otherwise, the LRDoS attack cannot have any effect on this system.We examine the impact of LRDoS attacks on a general feedback control system from two perspectives. First, we formally show in Proposition 1 that an LRDoS attack makes the steady-state error (i.e., es (t)) oscillate according to the attack pattern instead of converging to zero (Section III-A).Second, by modeling the system under an LRDoS attack as a switched system [19], [45], we prove the existence of LRDoS attacks that can force the victim system to diverge away from the desired state and constraint it to a state determined by the attacker through Proposition 2-3. To the best of our knowledge, we are the first to theoretically reveal how an LRDoS attack degrade the performance of feedback-control based systems.

The insights obtained in this section motivate us to propose a novel methodology to systematically analyze the impact of an LRDoS attack on a specific system. It includes four steps: (1) model the victim system under attack as a switched system; (2) quantify the impact on each subsystem and identify various attack scenarios; (3) determine the condition for an LRDoS attack to make the victim system Lyapunov stable and Lagrange stable; (4) decide the bound of the victim system's state and the relationship between the bound and the LRDoS attack. We apply this methodology to examine a web server .

A. Steady-State Error

An LRDoS attack transmits intermittent attack pulses to a target system. Let τk , k ∈ Z+ , represent the interval between the k t h and (k + 1)t h attack pulses, during which the LRDoS attack sends nothing. Note that the attack is not necessarily periodic (i.e., τk1 is not necessarily equal to τk2 , k1 = k2). For simplifying the ensuing discussion, we assume that the requests in each attack pulse have the same arrival rate, denoted as λa . Then, we can model the LRDoS attack as Consequently, we have d(t) =

λa + λn t = Tk .

We use capital letters to denote the Laplace transform of a component (i.e., G(s) and H (s)) and the input (i.e., D(s)), and employ to denote the convolution operator.
Although a feedback control system aims at minimizing the control error in the steady state, Proposition 1 shows that the steady-state error es (t)(= limt →∞ e(t)) cannot become zero (i.e., the system cannot stay at or converge to the desired state). More precisely, under an LRDoS attack, the control error oscillates with the attack period, and its magnitude is affected by λa . Hence, an attacker can cause different levels of damage by varying the attack period and λa .

Proposition 1: In the presence of an LRDoS attack, the steady-state error es (t), steady-state output ys (t), and steady-state control signal u s (t) oscillate according to the attack period τk , k   Z+ and λa .

B. Switched System Model

Let x(t) be the system state consisting of y(t) and u(t) [44]. A dynamic system is represented by x = f (x) and a˙solution of this equation corresponds to a curve in the state space that is also referred to as a system trajectory [44].We use x(t) = x z (t) + x n (t) and x (t) = x z (t) + x n (t) + x a (t) to denote system states in the absence/presence of an LRDoS attack, respectively. x z , x n, and x a are system states caused by zero input, normal requests, and attack requests, individually.

Compared with x(t), the extra component x a (t) in x (t) denotes discrete events in the system model f (x), because its components (i.e., ya (t) and u a (t)) appear at Tk , k   Z+ .A continuous-time system with discrete switching events is referred to as a switched system, which consists of a family of continuous-time subsystems and a switching rule that governs the switching between them [19], [45].

We first consider a family of subsystems given by x =˙ f p (x), where p   P, P is a finite index set, and for each p   P, f p is Lipschitz continuous. A switched system consists of a sequence of these subsystems

x = fσ (x),˙

where σ (t) : [0, ∞) → P is an index of the active subsystem. T When t   [tk , tk+1), σ (t) = i k , k, i   Z+ (i.e., the i k h subsystem is active in t   [tk , tk+1)).

Hence, σ (t) is a piecewise constant. At time instant tk+1 , σ (t) changes from i k to i k+1 , and therefore we call tk , k   Z+ , the switching times. The state x(t) of the switched system (1) is defined as the state x ik (t) of the i kh subsystem when t   [tk , tk+1). There are two types of switching points: time-dependent switching point and state-dependent switching point. The former is determined by attack pulses (i.e., switching happens at Tk). The latter is caused by the system (i.e., si , i   Z+).We use tk to denote all switching points (i.e., tk = {Tk , sk }).For ensuring the causality of the switching times (i.e.,tk+1 > tk > 0), we assume that if there are an infinite number of switching times, there exists ϑ > 0 such that for every constant T ≥ 0 one can find a positive integer k for which tk+1 − ϑ ≥ tk ≥ T .

IV. ATTACKING A WEB SERVER

Besides the qualitative analysis of the impact of an LRDoS attack on a general feedback control system, we further quantify it through real systems. As web servers are becoming the major platform for providing Internet services, we conducta comprehensive investigation into the impact of an LRDoS attack on the web server described in [16]. As another example, we investigate the impact of an LRDoS attack on an IBM Notes server proposed in [21] (Section V). It is worth noting that our methodology can be applied to other feedback-control based Internet services.For the web server, we address three challenging questions .

1. Are there other types of LRDoS attacks besides the one studied in ? If yes, what are they?

Section III reveals that there are other types of LRDoS attacks in addition to the one examined in , because the attack in allows the system to return to the steady state whereas Propositions 2-3 in Section III-C prove that an LRDoS attack can force the system to stay at an equilibrium point determined by the attacker. Motivated by this insight, we identify three types of LRDoS attacks in Section IV-B. It turns out that the attack in [16] is one type of LRDoS attacks, which is less severe than the other types in terms of the damage incurred to the web server.

2. If the answer to the first question is true, what are the impact and the effectiveness of these attacks?

AS Section IV-B uncovers new LRDoS attacks, we quantify their impact on the web server. It is non-trivial to model the impact of these new attacks compared to the one in ,because for the new attacks we have to determine the system state immediately before a new attack pulse arrives. For the attack in , the system is in the steady state before an attack general feedback control system, we prove in Proposition 4(Section IV-C) that a periodic LRDoS attack will cause the web server's state to converge with a periodic solution. Moreover, Proposition 5 gives closed-form equations for the maximal and minimal values of the admission rate constrained by different types of LRDoS attacks. Beside examining the impact, we also investigate the effectiveness of the LRDoS attacks by defining two metrics and modeling the relationship between the metrics and the parameters of an LRDoS attack (Proposition 6-7). The result implies the existence of optimal attack patterns, which will be studied in future work.

3. What kind of LRDoS attacks can make the web server under attack Lyapunov and Lagrange stable? What is the

ISSN 2394-3777 (Print)
ISSN 2394-3785 (Online)
Available online at www.ijartet.com

International Journal of Advanced Research Trends in Engineering and Technology (IJARTET)
Vol. 2, Issue 2, February 2015

bound of the state of the web server under attack?

A. The Web Server Model

the web server model. It employs a Proportional-Integral (PI) controller to adjust the admission rate (i.e., α(t)) according to the difference between the desired utilization (i.e., ρ) and the actual utilization (i.e., ρ(t)), which

is affected by the number of backlogged requests. Therefore, the system state can be described by the admission rate α(t),

the utilization ρ(t), and the number of backlogged requests n(t), as follows: where μW is the service rate and ρ(t) is a piecewise function with constants A, B, C, D, and υ. Note that this system is the same as the one in [16], except that we adopt a continuous-time model because it is more realistic and assume that μW is constant for analytical tractability.

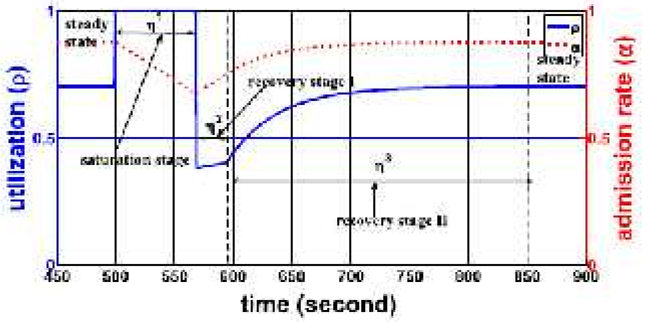

Fig. 3. The effect of one attack pulse at t = 500 on the admission rate and utilization.

Following [16], we assume that the arrival rate of normal requests is a constant λW . We evaluate the effects of stochastic arrival processes in Section VI-A and the supplementary material. Similar to Section III, we assume the arrival rate of attack requests in each attack pulse is a constant λW .We assume that the desired utilization ρ lies in the range of [Aυ + B, 1]. It is easy to extend the result to the scenario when ρ is between 0 and Aυ + B. Fig. 2(b) shows the relationship between ρ(t), ρ , and n(t), where 1 ≥ ρ ≥ Aυ + B.

B. Different Types of the LRDoS Attack

The goal of an LRDoS attack is to throttle the web server's admission rate so that requests from normal users are dropped.To achieve this, an attacker sends intermittent attack pulses to cause transient congestion in the web server, which forces the server to decrease its admission rate. On the arrival of an attack pulse, the server moves through three different stages before returning to the steady state: saturation, recovery I, and recovery II, as shown in Fig. 3. We use ηW ,1 , ηW ,2 , and ηW ,3 to denote the durations of these three stages. The saturation stage begins right after the arrival of an attack pulse. During this stage, the utilization equals 1 (i.e., ρ(t) = 1) and the admission rate decreases. After ρ(t) < ρ , the server enters two recovery stages consecutively,during which the admission rate and the utilization restore to the steady state. The difference between the two recovery stages lies in the model for ρ(t) and n(t) in Eqn. (3).Saturation stage: Once an attack pulse arrives, the system enters the saturation stage with ρ(t) = 1. The system state during this stage is characterized by ρ(t) = 1, α(t) = K (ρ − 1),˙ 1 W

−1)t 2 +(λW α −μW)t +(λW +λW)α ,and n(t) = 2 λn K (ρ00nan
where α0 is the initial value of α(t) and the expression for n(t) is backlogged requests is reduced to n(t) = (ρ − D)/C if

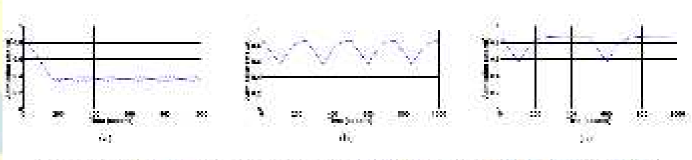

Fig. 4. Three types of LRDoS attacks and their impact on the admission rate. (a) Type I attack. (b) Type II attack. (c) Type III attack.

ρ ≥ Aυ + B, we have ρ(t) ≤ ρ and this stage ends. ηW ,1 can be obtained by solving n(ηW ,1) = (ρ − D)/C with the initial conditions [α0 , ρ0 , n 0], where α0 = α(0−) = α(0+), n + = (λW + λW)α0 + n − , and ρ0 = ρ(0+) = 1: at thena00 bottom of this page.

Recovery stage I:
At the beginning of this stage, because ρ(ηW ,1) ≤ ρ , α(t) stops decreasing and begins increasing. Consequently, ρ(t) also increases. The initial conditions for this stage include n(ηW ,1−) = n(ηW ,1+) = λW α(ηW ,1+),nα(ηW ,1−) = α(ηW ,1+) = α0 + K (ρ − 1)ηW ,1, and ρ(ηW ,1+) = AλW α(ηW ,1+)+ B. The evolution of the system state is given by ρ(t) = AλW α(t) + B, α(t) = K (ρ − ρ(t))˙n W α(t) − μW . This stage ends when n(t) = υ, and n(t) = λn˙and then the system enters the recovery stage II. ηW ,2 can be obtained by solving ρ(ηW ,2) = Aυ + B with the initial conditions α(ηW ,1+), ρ(ηW ,1+), and n(ηW ,1+):

Recovery stage II:
The differences between recovery stages I and II lie in the parameters and the initial conditions. The initial conditions here are α(ηW ,2−) = α(ηW ,2+) = υ ,ρ(ηW ,2−) = ρ(ηW ,2+) = Aυ + B, and n(ηW ,2−) =n(ηW ,2+) = λW α(ηW ,2+) − μW . This stage ends when the utilization reaches the desired value. Thus, ηW ,3 can be obtained by solving ρ(ηW ,3) = ρ with the initial conditions α(ηW ,2+), ρ(ηW ,2+), and n(ηW ,2):

According to the relationship between the attack period and the duration of the three stages, we identify three types of LRDoS attacks that have different impacts on the web server. Fig. 4 demonstrates the admission rate's trajectory in the presence of the different LRDoS attacks.

Type I attack: It has τk+1 < ηk ,1 + ηk ,2 . Under such an attack, the admission rate's trajectory involves two stages:

the saturation stage and the recovery stage I, as illustrated in

Fig. 4. When ηk ,1 < τk+1 < ηk ,1 + ηk ,2 , a new attack

pulse arrives during recovery stage I.When τk+1 ≤ ηk ,1 , a new attack pulse reaches during the saturation stage

The (k+1)t h attack pulse arrives when the system is still saturated due to the requests in the k t h attack pulse. In the extreme case, it becomes the flooding-based DoS attack..

C. The Impact of Periodic LRDoS Attacks

As it is easy for an attacker to launch a periodic LRDoS attack that has a constant interval between attack pulses, we analyze the impact of such attack on the web server and leave the theoretical investigation of non-periodic LRDoS attacks to future work. In Section VI and the supplementary material, we evaluate different non-periodic LRDoS attacks and compare them with the periodic LRDoS attacks through experiments.

We quantify the impact of a periodic LRDoS attack from three aspects:

- Proposition 4 proves that the victim system's state oscillates along with the attack.

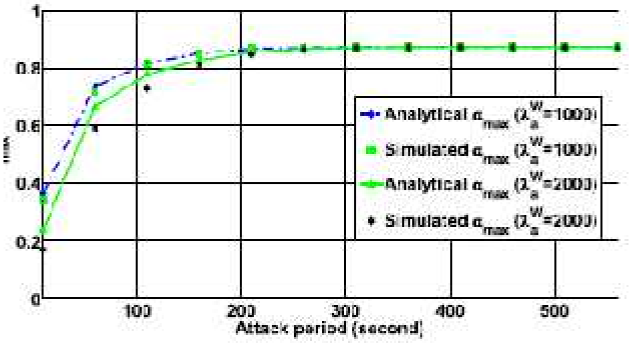

Fig. 5. α_{max} for different values of τ and λ_a^W.

E. Stability Analysis

We obtain the following results through stability analysis, including (1) the Lyapunov function (i.e., Eqn. (14)) for proving the system's stability; (2) the condition for an LRDoS attack to make the IBM Notes server Lyapunov stable (Proposition 12); (3) the relationship between the bound of the system trajectory and the parameters of an LRDoS attack (Proposition 13);We rewrite Eqn. (13) as follows: (1) u(t) = f 1 (u) = 0;˙ − m(t)); (3) u(t) = f (u) = 0.(2) u(t) = f 2 (u) = K i (m˙˙3) The roots of f p = 0, p {1, 2, 3} are the equilibriums points, including u e1 = 1, u e2 = u and u e3 = u . We ignore the case of u e1 because u(t) = 1 is a constant. To make the origin be the equilibrium point such that f p (0) = 0 for all p {1, 2, 3} [44], we let x(t) = u(t) − u e p for σ (t) = p, where u e p , p = {2} for the type I attack and p = {2, 3} for the type II attack, is the equilibrium point of the switched system Proposition 12 proves that an LRDoS attack with intervals larger than the saturation period (i.e., Tk+1 − Tk > ηk ,1 for all k) can make the system Lyapunov stable. In other words, such LRDoS attacks can force the system to stay away from the steady state. Otherwise, the attack becomes the flooding-based DoS attacks and u(t) converges to 1.After proving that the system under attack is Lagrange stable, Proposition 13 further establishes the relationship between the bound of the system state and the parameters of an LRDoS attack by proving u(t) is bounded .

VI. EXPERIMENTS

We carry out extensive experiments to evaluate the LRDoS attacks on the two servers. For the web server, we use the parameters from [16]: A = 0.00267, B = 0.2, C = 0.024,D = −1.4, υ = 75, K = 0.01, μ = 90, and ρ = 0.7. For the IBM Notes server, we adopt the parameters from : as = 0.6371, bs1 = 0.1692, bs2 = −0.1057 and K i = 0.1.The major reason for using the parameters from the original paper is that they make the system stable in the absence of LRDoS attacks. It is worth noting that while practical feedback control systems are necessarily stable [18] how to turn the controllers' parameters to make different systems stable is still an active research topic and out of the scope of this paper. In these experiments, we vary the parameters of the LRDoS attacks and the input of legitimate users to the servers, and examine the corresponding impact. Section VI-A and Section VI-B present the Mat lab simulation result of the web server and the IBM Notes server, respectively. Due to the page limit, we leave the test bed experiment results and many other simulation results in the supplementary material.

A. Simulation Result for the Web Server Model

In Section IV, we analyze the effect of the attack period τ on φ W (i.e., percentage of normal requests dropped) and ψ W (i.e., number of normal requests dropped per attack request). We evaluate such effect by launching LRDoS attacks with λW = 1000, 2000 requests per second and a wide range of attack periods. Note that the λW s are high enough to force the web server to enter the saturation stage on the arrival of each attack pulse. We also examine small λW s that cannot saturates the web server by individual attack pulse, and report the results in the supplementary material. In the experiments, the attacker sends either periodic pulses or random pulses. For the latter, the interval between consecutive attack pulses follows exponential, normal and Pareto distributions. We only show the results of the exponential distribution, and leave the others in the supplementary material. Moreover, we simulate both constant and Poisson arrival rate for normal requests. φ W decreases with τ , because the web server has longer time to recover its admission rate and consequently takes in more normal requests. Hence, if an attacker wants to cause more legitimate requests to be dropped, she should use a smaller attack period. The extreme case is the flooding attack, whose period is zero. We also observe that φ W increases with λW, because a larger λW causes severe damage. However, when we consider the attack cost (e.g., the number of attack request), a shorter interval or a larger λWa may not be preferred. Fig. 12 shows that ψ W increases with τ. That is, a larger attack period yields more damage per attack request. Moreover, ψ W converges to a constant as τ increases, because in this situation the attacks belong to the type III attack and the web server can return to the steady state during the intervals between consecutive attack pluses. In other words, enlarging τ will not increase the number of dropped requests. Similarly, a larger λW may not be cost-effective as Fig. 12a W decreases with λW shows that ψa

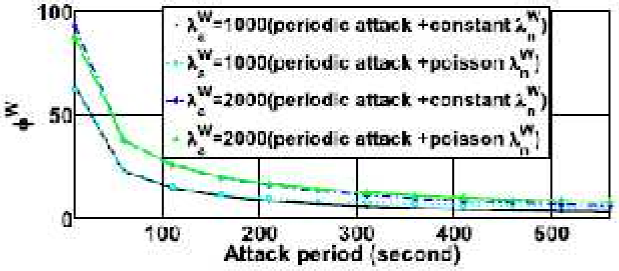

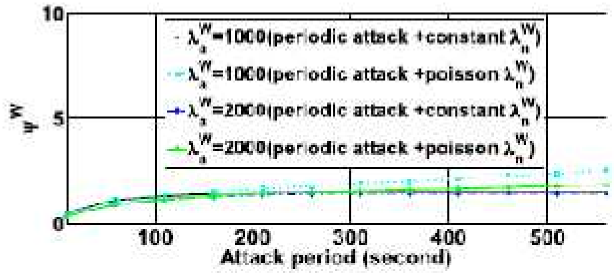

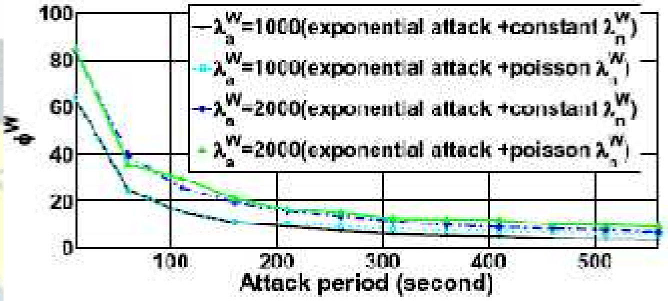

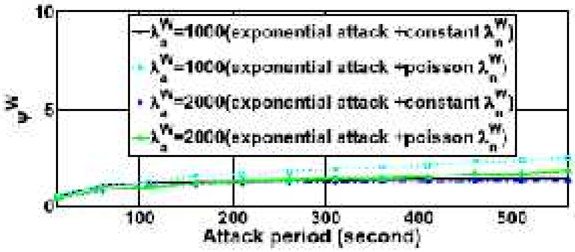

the differences in the effectiveness of LRDoS attacks due to the different arrival processes of normal requests.

We can observe from that both φ W and ψ W under the constant-rate process are less than those under the Poisson process. The reason may be that the bursts in the Poisson process were superposed on attack pulses and then caused more normal requests to be dropped, not to mention majority of dropped requests that the large bursts alone can also cause request losses in the absence of attack. The difference for ψ W under two arrival processes increases with τ . It may result from the different types of the attack. More precisely, when τ increases, the attack evolves from type I, type II to type III. Under the type III attacks, the total number of dropped normal requests caused by each attack pulse is fixed if the arrival rate of normal request remains constant. In contrast, the bursts in the Poisson arrival processes along with the attack pulses may force more normal requests to be discarded. Under the type I and II attacks (i.e., τ is small), the difference caused by different arrival processes is not obvious, because the attack pulses lead to the Poisson process) becomes more obvious when τ increases similar results in terms of the relationship between φ W (or ψ W) with τ when the intervals between consecutive attack pulses follow the exponential distribution. However, comparing Fig. 13 and Fig. 11, we find that using the same λW the LRDoS attack with randomized interval causes less damage than the attack with fixed interval, especially for small τ . The reason may be that on one hand when the interval is longer than the fixed attack period the attack may allow the server to have more time to recover. On the other hand, when the interval is shorter than the fixed attack period, the current attack pulse may cause the server to drop more attack requests in the next attack pulse, thus moderating the attack damage. Moreover, similar to Fig. 12,Fig. 14 shows that the difference for ψ W under two arrival processes of normal requests (i.e., constant-rate process versus).

VII. Discussion

The goal of this paper is to reveal the vulnerability of feedback-control based systems to the LRDoS attacks through theoretical analysis and then propose a new methodology quantify the impact of the LRDoS attacks on such systems. Therefore, we assume that the feedback control model for the victim system is available, such as, the web server model in and the IBM notes server model . However, some feedback-control based systems may not have constructed the model. For example, the system in just measures the output (i.e., y(t)) of the process (i.e., h(t) without giving the equations of h(t). In order to apply our methodology to analyze such systems, users can first model h(t) through system identification . There are also systems that do not use feedback controller or employ other kinds of controllers, such as adaptive control, model predictive control, robust control, etc. Our methodology could not investigate the impact of LRDoS attacks on such systems and we will examine them in future work. To simplify the theoretical analysis, we assume that the arrival rates of normal requests to the two servers are constants. Although this assumption may not be realistic, our analysis sheds light on the impact of the LRDoS attacks. Other arrival processes along with the attack may cause severer damage. The reason is that the bursts in the arrival process of normal users may be superposed on the attack pulses and consequently cause more damage to the server, not to mention that the large bursts alone may also affect the server. The experiment results in Section VI demonstrate it. In future work, we will enhance our model by considering more realistic arrival process models.

VIII. Conclusion

We investigate the vulnerability of feedback-control based Internet services to the LRDoS attacks. We first examine the impact of the LRDoS attacks on a general feedback control system and prove that LRDoS attacks can force the system's steady-state error to oscillate along with the attack. By modeling the system under attack as a switched system, we prove the existence of LRDoS attacks that can drive the system to a state other than the desired state. Both the oscillation of steady-state error and staying away from the desired state impair the system's performance. Then, we propose a novel methodology to analyze the impact of LRDoS attacks on specific feedback control systems. We obtain many new insights by applying the methodology to examine a web server model and an IBM Notes server model. In future work, we will investigate the tradeoff between the effectiveness and the cost of LRDoS attack, and design defense mechanism to mitigate the damage of LRDoS attack.

References

[1] T. Abdelzaher, K. Shin, and N. Bhatti, "Performance guarantees for web server end-systems: A control-theoretical approach," IEEE Trans. Parallel Distrib. Syst., vol. 13, no. 1, pp. 80–96, Jan. 2002.

[2] J. Hellerstein, Y. Diao, S. Parekh, and D. Tilbury, Feedback Control of Computing Systems. Hoboken, NJ, USA: Wiley, 2004.

[3] M. Welsh and D. Culler, "Adaptive overload control for busy internet servers," in Proc. USENIX Symp. Internet Technol. Syst., 2003, pp. 1–4.

[4] Y. Lu, T. Abdelzaher, C. Lu, L. Sha, and X. Liu, "Feedback control with queueing-theoretic prediction for relative delay guarantees in web servers," in Proc. 19th IEEE RTAS, May 2003, pp. 208–217.

[5] H. Lim, S. Babu, J. Chase, and S. Parekh, "Automated control in cloud computing: Challenges and opportunities," in Proc. 1st Workshop ACDC, Jun. 2009, pp. 13–18.

6] Y. Seung, T. Lam, L. Li, and T. Woo, "CloudFlex: Seamless scaling of enterprise applications into the cloud," in Proc. IEEE INFOCOM, Apr. 2011, pp. 211–215.

[7] A. Sharifi, S. Srikantaiah, A. Mishra, M. Kandemir, and C. Das, "METE: Meeting end-to-end QoS in multicores through system-wide resource management," ACM SIGMETRICS Perform. Eval. Rev., vol. 39, no. 1 pp. 13–24, Jun. 2011.

[8] T. Abdelzaher, J. Stankovic, C. Lu, R. Zhang, and Y. Lu, "Feedback performance control in software services," IEEE Control Syst., vol. 23, no. 3, pp. 74–90, Jun. 2003.

[9] S. Park and M. Humphrey, "Predictable high-performance computing using feedback control and admission control," IEEE Trans. Parallel Distrib. Syst., vol. 22, no. 3, pp. 396–411, Mar. 2011.

[10] Z. Wang, Y. Chen, D. Gmach, S. Singhal, B. Watson, W. Rivera, et al.

"AppRAISE: Application-level performance management in virtualized server environments," IEEE Trans. Netw. Service Manag., vol. 6, no. 4, pp. 240–254, Dec. 2009.

[11] K. Kim and P. Kumar, "Cyber–physical systems: A perspective at the centennial," Proc. IEEE, vol. 100, no. 13, pp. 1287–1308, May 2012.

[12] M. Huebscher and J. McCann, "A survey of autonomic computing—Degrees, models, and applications," ACM Comput. Surv., vol. 40, no. 3,pp. 1–7, Aug. 2008.

[13] A. Kuzmanovic and E. Knightly, "Low-rate TCP-targeted Denial-of-service attacks: The shrew vs. the mice and elephants," in Proc.Conf. Appl., Technol., Archit., Protocols Comput. Commun., Aug. 2003, pp. 75–86.

[14] M. Guirguis, A. Bestavros, I. Matta, and Y. Zhang, "Exploiting the transients of adaptation for RoQ attacks on internet resources," in Proc 12th IEEE ICNP, Oct. 2004, pp. 184–195.

[15] X. Luo, R. Chang, and E. Chan, "Performance analysis of TCP/AQM under Denial-of-service attacks," in Proc. 13th IEEE Int. Symp.MASCOTS, Sep. 2005, pp. 97–104.

[16] M. Guirguis, A. Bestavros, I. Matta, and Y. Zhang, "Reduction of quality (RoQ) attacks on internet end-systems," in Proc. IEEE 24th Annu. Joint Conf. Comput. Commun. Soc., vol. 2. Mar. 2005, pp. 1362–1372.

[17] M. Guirguis, A. Bestavros, I. Matta, and Y. Zhang, "Adversarial exploits of end-systems adaptation dynamics," J. Parallel Distrib. Comput., vol. 67, no. 3, pp. 318–335, 2007.

[18] K. Ogata, Modern Control Engineering, 5th ed. Englewood Cliffs, NJ, USA: Prentice-Hall, 2009.

[19] J. Lunze and F. Lamnabhi-Lagarrigue, Handbook of Hybrid Systems Control: Theory, Tools, Applications. Cambridge, U.K.: Cambridge Univ. Press, 2009.