

Exploiting the Transients of Adaptation for RoQ Attacks on Internet Resources


Azer Bestavros

Joint work with
Mina Guirguis & Ibrahim Matta



Denial of Service (DoS) Attacks

- How: Subject a service to a load that exceeds its steady-state capacity...



- Goal: Make resource unavailable to legitimate users...

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources

2

Denial of Service Attacks

- Most Recent Example: Attack on SCO's Web site on 2/2/04 courtesy of MyDoom
 - A nuisance?
 - Freedom of Expression?
 - Act of Patriotism?
 - How about \$26.1B of lost productivity!!!

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources

3

DoS: The Good News

- Need lots of resources to mount attack
 - Attack can be anticipated
 - *There is a Red Sox Game; stay home!*
- Easy to detect intrusion/attack
 - Attack signature is simple
 - *When was the last time you saw elephants crossing the BU Bridge?*
- Theoretically can trace back perpetrators
 - *If nothing else, getting a ticket is a deterrent – maybe ☺*

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources

4

What If...

- It takes less resources to mount attack
 - Attack cannot be anticipated
 - *It's just another day of gridlock on the BU Bridge*
 - Intrusion/attack detection is hard
 - Attack signature is not as simple!
 - *No suspicious traffic patterns seen on bridge*
 - It is hard to trace back perpetrators
 - Attacker is keen on not being exposed (a.k.a. "Reduction of Liability", or RoL)
- Say hello to the world of "low-rate" DoS attacks

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources

5

Prior Work: Shrew Attack [KK03]

- Goal: Cause a subset of TCP flows to perpetually timeout
- How: Synchronize the attack traffic to exploit TCP timeout mechanism
 - Upon a loss of a window of packets, a TCP flow is forced into timeout
 - minRTO is lower bounded by 1 sec [RFC2988]

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources

6

Our Work: RoQ Attacks



- **Goal:** “Bleed” the system of its capacity by forcing it to operate in its most inefficient region—with *minimal exposure*
- **How:** “Exploit” built-in load adaptation mechanisms to make the system perpetually in a transient state—unstable
- **Hint:** Make other drivers brake when they should accelerate and accelerate when they should brake. Just be a Boston driver ☺

Our Work: RoQ Attacks



	What it Exploits	What it targets	What it needs	Intrusion detection	Trace-back
DoS	Steady-state	Any system Web Servers: DNS: Internet Routers	Lots Find Elephants and herd them to bridge	Easy Watch for elephants!	Easy Find elephant owners!
Shrew	TCP Timeout Mechanism	TCP flows using timeout mechanisms	Depends on targeted victims' RTT	Easy Elephants spotted periodically	Hard Spoofing can be used
RoQ	Adaptation Dynamics	Any adaptation TCP/AQM; BGP; Admission Control; Load Balancers; sensor coordination	Few Attack goal is to maximize damage while minimizing exposure	Hard Transients could occur under normal operation; gridlock on the bridge	Harder Spoofing can be used for both source and destination

Adversarial Exploits of Adaptation



- Adaptation mechanisms are built under an assumption of a non-adversarial loads
 - Examples: random traffic patterns, random arrival processes, etc.
- Questions:
 - What load patterns would be most virulent to a given adaptation scheme?
 - How much adversarial load would it take to make adaptation harmful?
 - ...

RoQ Attack: Definition



- RoQ Attacks maximize the marginal utility of attack traffic → Potency

$$\Pi = \frac{\text{Damage}}{\text{Cost}^\dagger}$$

- Many Possible Instantiations
 - Damage = Rejected requests, response time, wasted BW
 - Cost = Injected requests, # of attackers, attack BW
 - Aggressiveness = Tolerance to exposure
 - Large Omega → Largest level of aggression
 - Small Omega → Minimal exposure

Adversarial Exploits of Adaptation

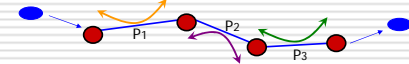


- We considered three examples of RoQ:
 - Congestion Control in Networks
 - Admission Control for Web Servers
 - Load Balancing in CDNs
- Many other vulnerabilities exist
 - Dynamic routing (e.g., BGP)
 - Power conservation in sensor networks
- Hard to find systems that would be safe!

RoQ: Fundamental Concept



- Adopted Kelly's optimization framework
 - Maximize users' utilities subject to the network's capacity constraints [K99]
 - Primal source algorithm



$$\frac{d}{dt} x_r(t) = \kappa \left(\underbrace{w_r - x_r(t)}_{\text{Additive Increase}} \underbrace{\sum_{l \in \mathcal{L}_r} p_l \left(\sum_{s \in \mathcal{L}_s} x_s(t) \right)}_{\text{Multiplicative Decrease}} \right)$$

RoQ: Fundamental Concept

- Link price functions reflect prices fed back to sources as the load on the links varies

Sources' algorithms iterate over rates
Links' algorithms iterate over prices

- Convergence and stability can be proved through Lyapunov function [K99]

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 13

RoQ: Fundamental Concept

- RoQ attacks will hinder convergence

During attack
Without attack

- Can destroy the "contractive mapping" of the pricing function

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 14

RoQ: Fundamental Concept

- Potency can be computed as a function of the convergence

$$\text{Potency} = \Pi = \frac{\text{Damage}}{\text{Cost}^{\delta}} = \frac{\delta(\frac{1}{x} + \frac{1}{y})}{(\delta/(\frac{1}{x} + \frac{1}{y}))^{1/\delta}}$$

- Example

- Optimal attack exists that maximizes the potency

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 15

Network Adaptation: Part Un

- A packet loss = congestion signal

- AIMD Control
 - No packet loss → increase sending rate linearly
 - Packet loss → decrease sending rate exponentially
- Timeout
 - Nothing is going through → shut off for exponentially longer periods of time

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 16

Network Adaptation: Part Deux

- What generates a loss?
 - No space in router queue (a.k.a. DropTail)
 - Drop packet if queue builds up (a.k.a. RED)
- RED as example of Active Queue Mgmt
 - Tries to avoid "herding behavior" by randomizing packet losses across flows and by relating loss probability to queue length

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 17

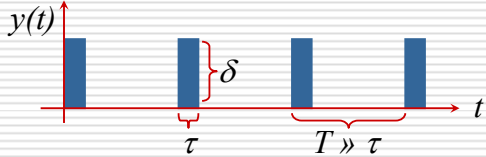
Premise of a RoQ & RoL Attack

- RoQ:** Attacker sends packets at high rate—enough to cause lots of flows to slow down exponentially fast (e.g., by halving their sending rate)
- RoL:** Attacker shuts off
- Resource will be underutilized until flows "rev-up" their sending rate, which is a slow linear process by design
- Go back to 1...

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 18

Attack Pattern

- A simple “square wave”

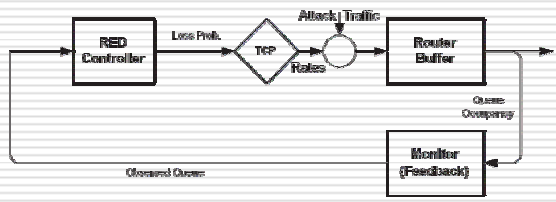


- Values of δ , τ , and T will depend on setting—stay tuned...

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 19

Network Adaptation: RED+TCP

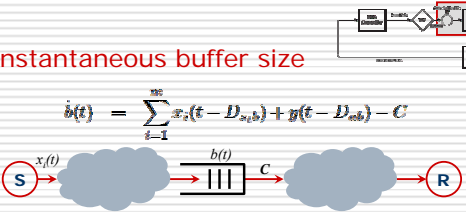
- Modeled as a set of difference equations for a set of n flows, each subject to a feedback control loop



Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 20

Network Adaptation: Router

- Instantaneous buffer size

$$\dot{b}(t) = \sum_{i=1}^{nq} x_i(t - D_{s_i,b}) + b(t - D_{cb}) - C$$


$$r_i(t) = D_i + \frac{b(t)}{C}$$

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 21

Network Adaptation: RED

- Queue (EWM) Average Size

$$\dot{v}(t) = -\beta C(v(t) - b(t)), \quad 0 < \beta < 1$$

- Loss Rate

$$p_i(t) = \begin{cases} 0 & v(t) \leq B_{min} \\ \sigma(v(t) - s) & B_{min} < v(t) < B_{max} \\ 1 & v(t) \geq B_{max} \end{cases}$$

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 22

Network Adaptation: TCP

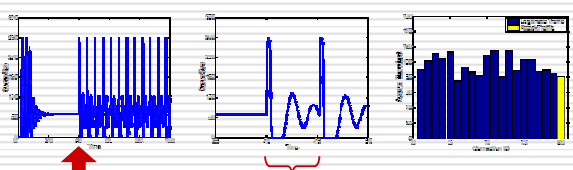
- Sending Rate Evolution (*a la* AIMD)

$$\dot{x}_i(t) = \frac{x_i(t - r_i(t))}{r_i^2(t) c_i(t)} (1 - p_i(t - D_{s_i,c_i})) - \frac{x_i(t) x_i(t - r_i(t))}{2} (p_i(t - D_{s_i,c_i}))$$

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 23



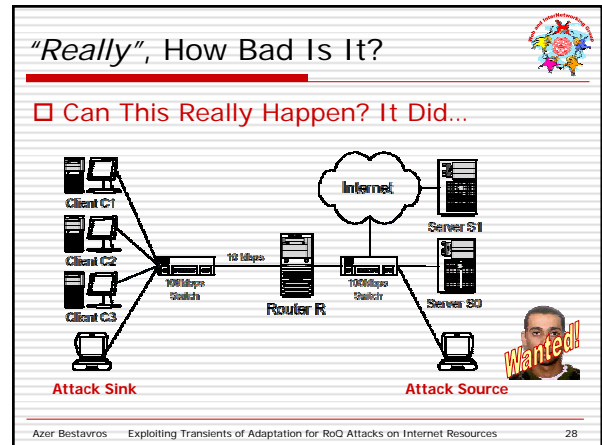
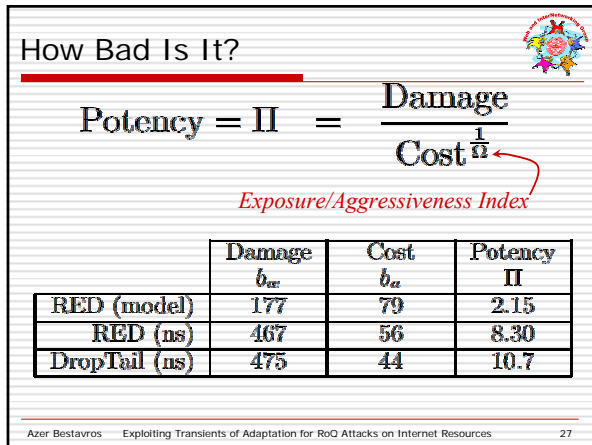
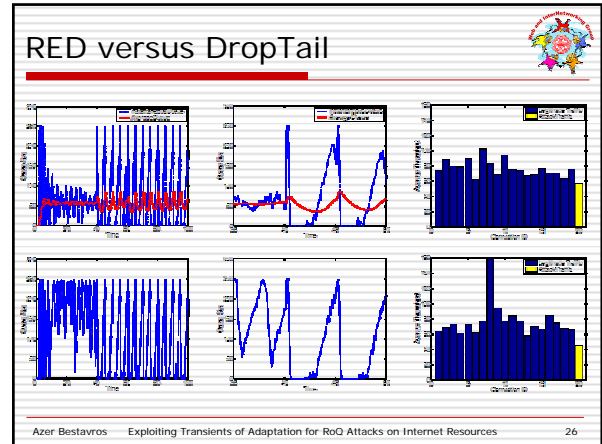
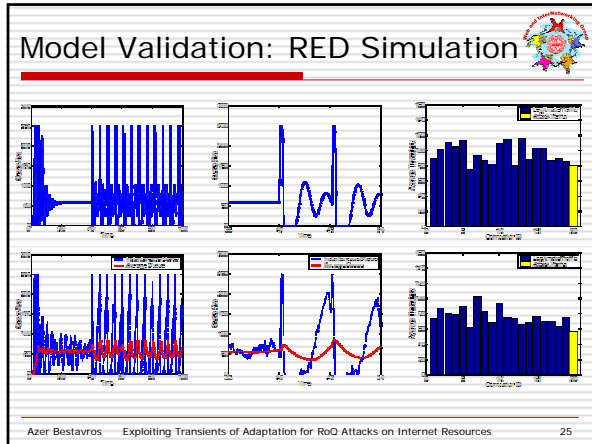
Network Adaptation: RED+TCP

- Model can be instantiated and numerical results obtained



Attack Starts Here Attack Period

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 24



It is even scarier...

- Could be mounted as a distributed RoQ
 - Time multiplexed
 - Traffic multiplexing
- Traceback is that much harder!
 - Spoofing source addresses
 - Attack sink does not even have to exist!

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 31

Stealing Quality (vs bandwidth)

$$\text{Potency} = \Pi = \frac{\text{Damage}}{\text{Cost}^{\frac{1}{n}}}$$

	Delay Jitter Before	Delay Jitter After	Damage (msec)	Cost b_n	Potency Π
RED (model)	0.0	28.5	28.5	79	0.36
RED (ms)	8.50	37.5	29.0	56	0.52
DropTail (ms)	32.0	42.0	10.0	44	0.23

... or how to make a RED link look like a DropTail link?

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 32

Tuning Attack Parameters

Could be done using an on-line controller

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 33

Adversarial Exploits of Adaptation

- We considered three examples of RoQ:
 - Congestion Control in Networks
 - Admission Control for Web Servers
 - Load Balancing in CDNs
- Many other vulnerabilities exist
 - Dynamic routing (e.g., BGP)
 - Power conservation in sensor networks
- Hard to find systems that would be safe!

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 34

Admission Control

- A "Gate" used to protect from overload
 - Admit (cross the bridge ☺)
 - Reject (into the river ☹)
 - Postpone

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 35

Admission Control Adaptation

- Admission Controller
 - What percentage of requests should be admitted?
 - Calculated based on the deviation between the server's state and a target value
 - PI Controller, AIMD Controller, etc...
- Feedback Monitor
 - Measures the server's state and report it back to the Controller
 - Feedback delay

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 36

Admission Control Adaptation

□ Modeled as a discrete time system

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 37

Admission Control: Model

□ Controller: Proportional Integrative (PI)

$$\alpha(i) = K \times (\underbrace{\mu^* - \rho(i)}_{\text{Error Signal}}) + \alpha(i-1)$$

□ Gate

$$m(i) = \alpha(i) \times \lambda(i)$$

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 38

Server: Model

□ Pending requests

$$n(i) = n(i-1) + m(i) - \mu(i-1)$$

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 39

RoQ Attack Premise

1. **RoQ**: Attacker sends requests at high rate in a very small period of time, enough to push the server into overload
2. **RoL**: Attacker shuts off
3. Admission control will shut off subsequent legitimate requests. Since the system is thrashing, recovery will take a longer time
4. Go back to 1...

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 40

RoQ Attack Pattern

□ A simple "square wave"

$$y(t) = \begin{cases} \delta & t \bmod T \leq \tau \\ 0 & \text{otherwise} \end{cases}$$

Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 41

Admission Control Adaptation

$$\text{Potency} = \Pi = \frac{\text{Damage}}{\text{Cost}^{\frac{1}{\alpha}}} \rightarrow \frac{R_j}{\delta \tau}$$

□ Model can be instantiated and solved

□ Large potencies possible "theoretically"


Azer Bestavros Exploiting Transients of Adaptation for RoQ Attacks on Internet Resources 42

Implementation Setup



Server: Minihttpd

- Admission Control
- Forks a cgi script
 - Access 1MB
 - ~ 20 msec




Clients: Httpperf

Utilization

- Memory utilization = Used / Total

Implementation Results



Attack Start at 120 (, 740, ...) with 800 requests; system recovers only at time 500 (, 1120, ...)


Potency depends on controller settings (e.g., gain) and other system characteristics

Implementation Results



Effect of Feedback Delay

- Real Delay
- Averaging (EWMA)



Limitations

- Linux alleviates thrashing: It Kills threads making collecting of data real hard
 - *Only able to cause moderate thrashing*
- Limitation on number of open connections generated by Httpperf
 - *Only able to use 4 machines*

Take Home Messages



RoQ Attacks Exploit Dynamics:

It is NOT capitalizing on a static property of a protocol—unlike the “shrew” attack which causes perpetual timeouts

RoQ Attacks Trade off Damage and Cost:

It is NOT aiming to take a resource down at any cost, but rather it is aiming to get the maximum damage per attack byte

Food For Thought...



More elaborate attacks

- Complex attack patterns

Are there fundamental tradeoffs?

- RoQ tolerance versus utilization/delay/fairness

Other adaptation susceptibilities

- Routing algorithms, sensor nets, ...

Countermeasures

- Randomized adaptation
- Intrusion detection
- Traceback

Exploiting the Transients of Adaptation for RoQ Attacks on Internet Resources

Mina Guirguis, Azer Bestavros, and Ibrahim Matta. *Exploiting the Transients of Adaptation for RoQ Attacks on Internet Resources*. In *IEEE ICNP 2004*.

Mina Guirguis, Azer Bestavros, Ibrahim Matta, and Yuting Zhang. *Reduction of Quality (RoQ) Attacks on Internet End-Systems*. In *IEEE INFOCOM 2005*.



<http://www.cs.bu.edu/groups/wing>

More information available from WING Publications
<http://www.cs.bu.edu/groups/wing>