Disabling a Computer by Exploiting Softphone Vulnerabilities

Ryan Farley and Xinyuan Wang
George Mason University
September 26, 2013

Threat and Mitigation

- **Introduction**
- Background
- Disabling the Softphone Host
- Defenses
- Experiments
- Conclusion

Introduction

- Many VoIP exploits stem from underlying SIP
  - De facto signaling protocol
- Previous works demonstrate protocol attacks
  - Remote monitoring, billing fraud, voice pharming
- Focus here is on the system hosting a softphone
  - Stability, security
  - Exploitable softphone in experiments is Vonage client
- And how to mitigate such threats

Specifically

- Two attacks that remotely disable host until reboot
  - A faster noisy attack effective in minutes
  - A slower but stealthier attack
- Two rapidly deployable defenses
  - Do not interfere with standard SIP operation
  - Threshold filtering inhibits arrival rate spikes
  - Limited Context Aware (LCA) filtering blocks only attack signals even at low arrival rates

- Introduction
- **Background**
  - Fundamental Problem
  - Invite Flooding
- Disabling the Softphone Host
- Defense
- Experiments
- Conclusion

Background

- Session Initiation Protocol (SIP)
  - Manages multimedia sessions
  - Between endpoints called User Agents (UAs)
  - Request-response paradigm
- Making a call
  - A sends an `Invite` to B
  - B's proxy sends a `100 Trying` back to A
  - B sends a `180 Ringing` back to A
  - If answered, B sends a `200 OK` to A, who Acks back

The SIP Behind a VoIP Call

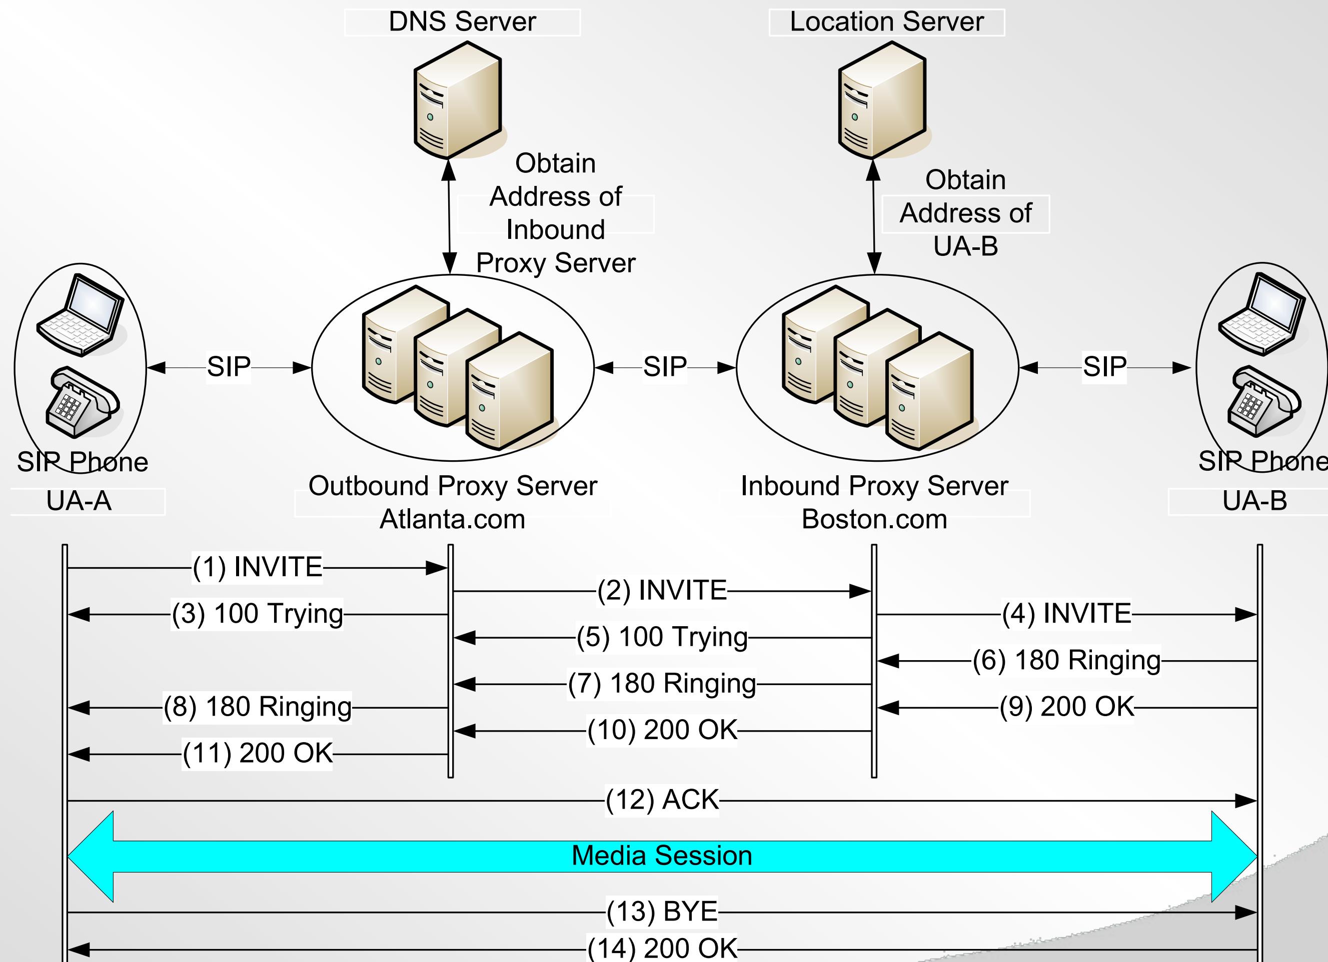

GEORGE MASON UNIVERSITY

Where Innovation Is Tradition

Fundamental Problem

- Invites are easy to spoof
  - Well known Invite flooding attacks
- SIP RFC provides for HTTP digest authentication
  - Invite, Register, Bye
  - From UAC to UAS, not required the other way around
  - Previous work shows Vonage, AT&T vulnerable
- Not nearly as widely implemented as it should be

Flooded Behavior

- Unattended softphone will ring until timeout
  - Will not ring for duplicate Call-IDs repeated within 60s
- Once all RTP ports reserved responds with Busy
  - Two ports mean two simultaneous ringing lines
  - Roughly only two spoofed Invites every 3 minutes needed to disrupt incoming calls
- Race condition inhibits outgoing calls

- Introduction
- Background
- **Disabling the Softphone Host**
  - Noisy Attack
  - Stealthy Attack
- Defense
- Experiments
- Conclusion

Disabling the Softphone Host

- Previous work targets infrastructure or devices
  - Not clear precisely how softphone weaknesses open host up for attack
- Two attacks
  - Can disable Windows XP machines running official Vonage softphone
  - First consumes memory resources in minutes
  - Second is slower but much stealthier

GEORGE MASON UNIVERSITY

Noisy Attack

- Memory allocated for every Call-ID seen
  - e.g., RFC requires 3 Busy signaling attempts over 10 seconds
  - Poor memory management impacts host
- Invite flood
  - Hundreds per second
  - Only need unique Call-ID

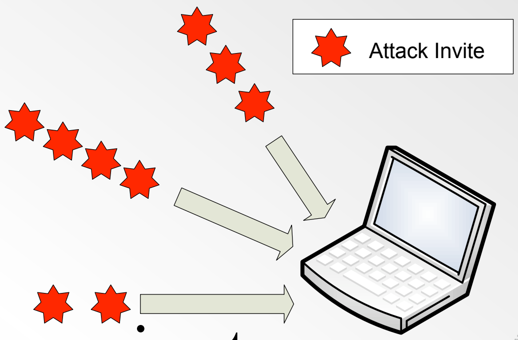

- Host begins to thrash within a few minutes
  - UI frozen at 16 minutes; unusable until reboot

Stealthy Attack

- Noisy, is well, noisy
  - Cancels can stop the ringing
  - Tells receiver to ignore Invites with same Call-ID
  - But memory consumption still happens
- Multiple Cancels
  - Secure chance of silence
  - Reduce arrival rate to 1/(n+1), with n cancels
- Same result, longer period, stealthier
  - Two hours

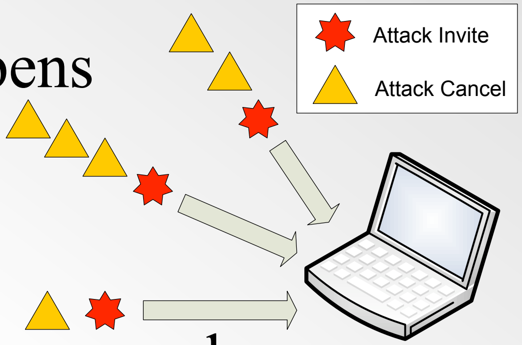

GEORGE MASON UNIVERSITY

- Introduction
- Background
- Disabling the Softphone Host
- **Defense**
  - Threshold
  - Limited Context Aware
- Experiments
- Conclusion

Defenses

- Must defend against single packet attacks
  - Group packets to be analyzed
- External factors help define meaningful calls
  - More than 1-2 calls a second beyond human threshold
    - Our first defense limits the rate of invites
    - But the second attack defeats this with its low arrival rate
  - If canceled unreasonably fast, then why ring at all?
    - Our second defense builds a context to stop meaningless calls

Threshold Filter

- Noisy attack makes finding signature difficult
  - Both in network and application layer
  - Only an arrival rate threshold indicates possible attack
- Some attack packets may pass, but very low rate
  - Phone would ring extended time, most likely alert user

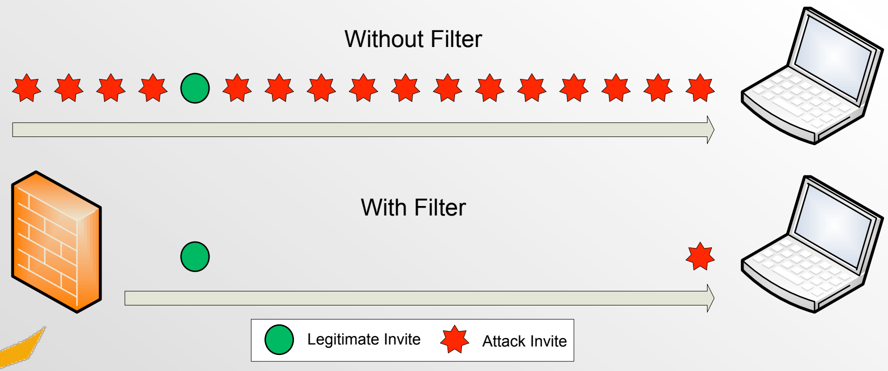

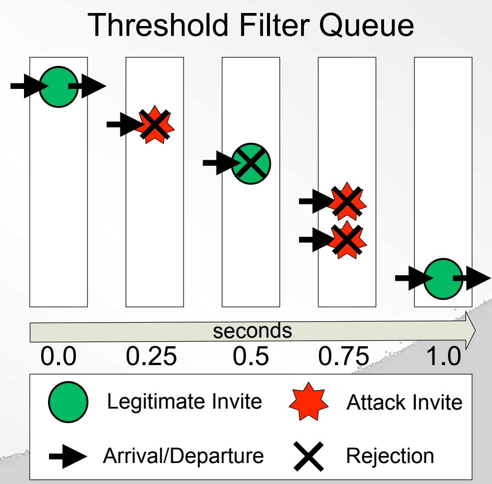

GEORGE MASON UNIVERSITY

Where Innovation Is Tradition

Limited Context Aware Filter

- Stealthy arrival rate is lower than noisy
  - Threshold filter not as effective
  - Signature: at least one Cancel per Invite
- Queue forms a limited, by time, context
  - Time is the acceptable delay to begin ringing
  - Determine if in that time any Cancels appear

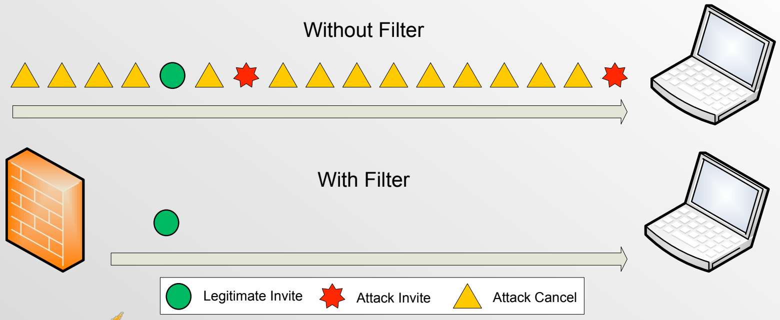

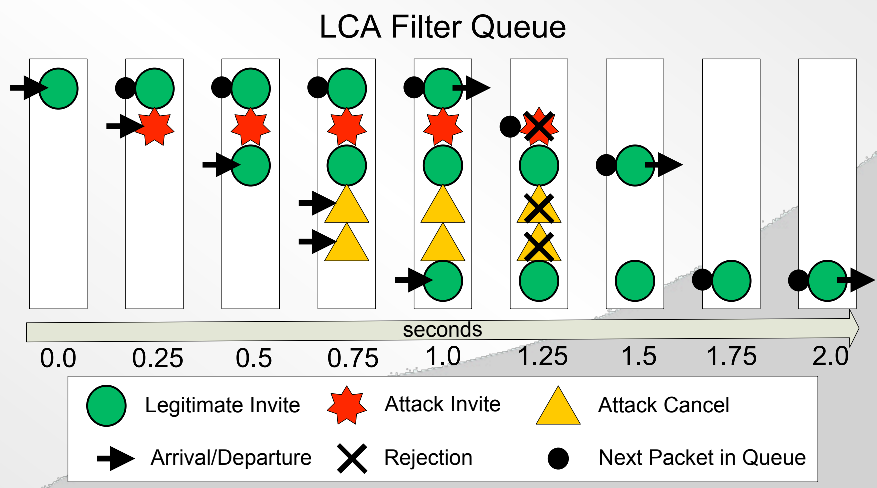

GEORGE MASON UNIVERSITY

Where Innovation Is Tradition

- Introduction
- Background
- Disabling the Softphone Host
- Defense
- **Experiments**
  - Attacks
  - Defense
- Conclusion

GEORGE MASON UNIVERSITY

Experiments

- Implementation
  - Attacks from Linux socket programs
    - Invite template from PCAP trace of legitimate call to target
  - Filters through FreeBSD divert sockets
    - Within a transparent network bridge
  - Targets were Windows XP virtual machines
    - 256 MB RAM
    - X-PRO Vonage 2.0 Softphone, release 1105x build 17305
  - Any unnecessary outbound traffic blocked at network's public edge to protect Vonage servers

Before Attack

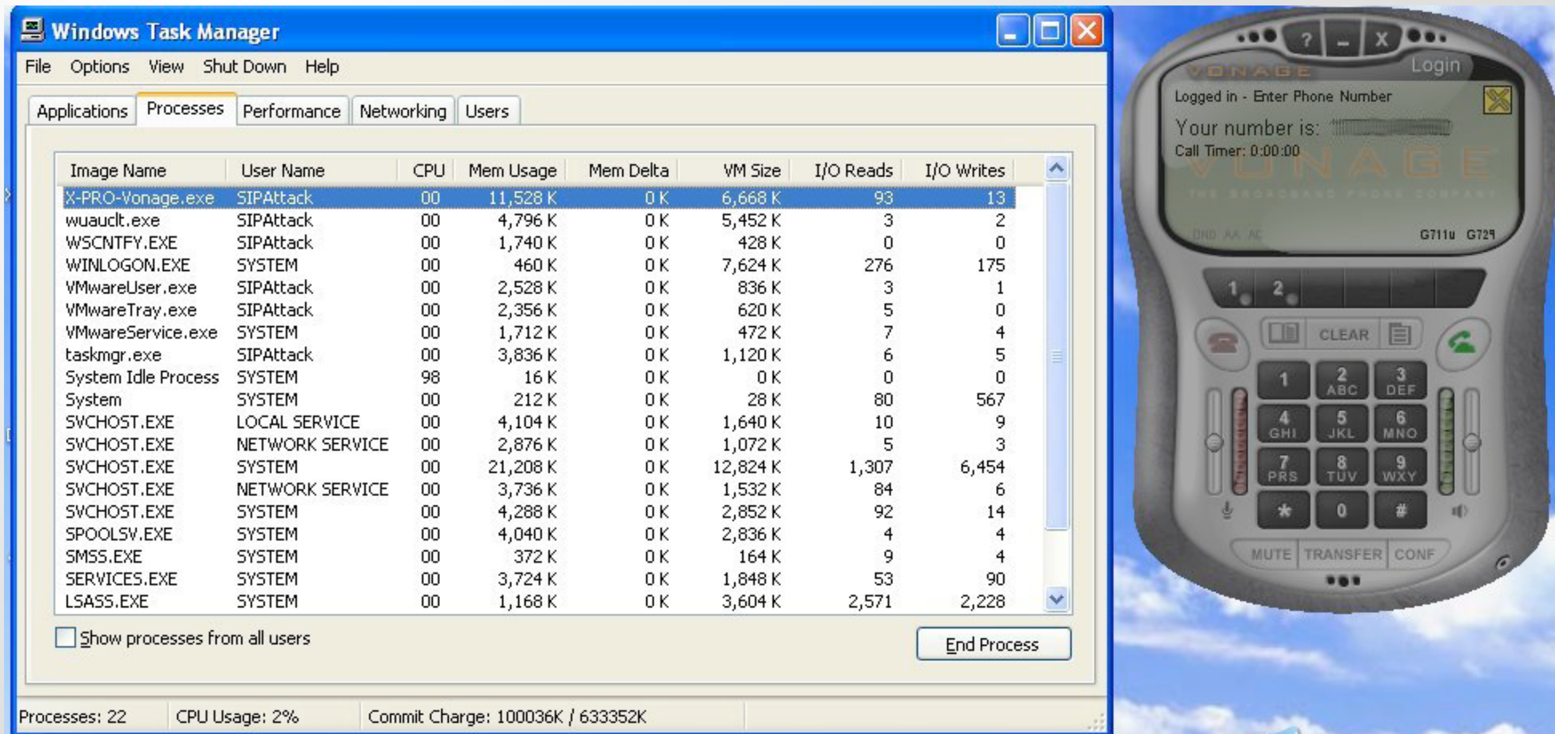

GEORGE MASON UNIVERSITY

Where Innovation Is Tradition

After Attack

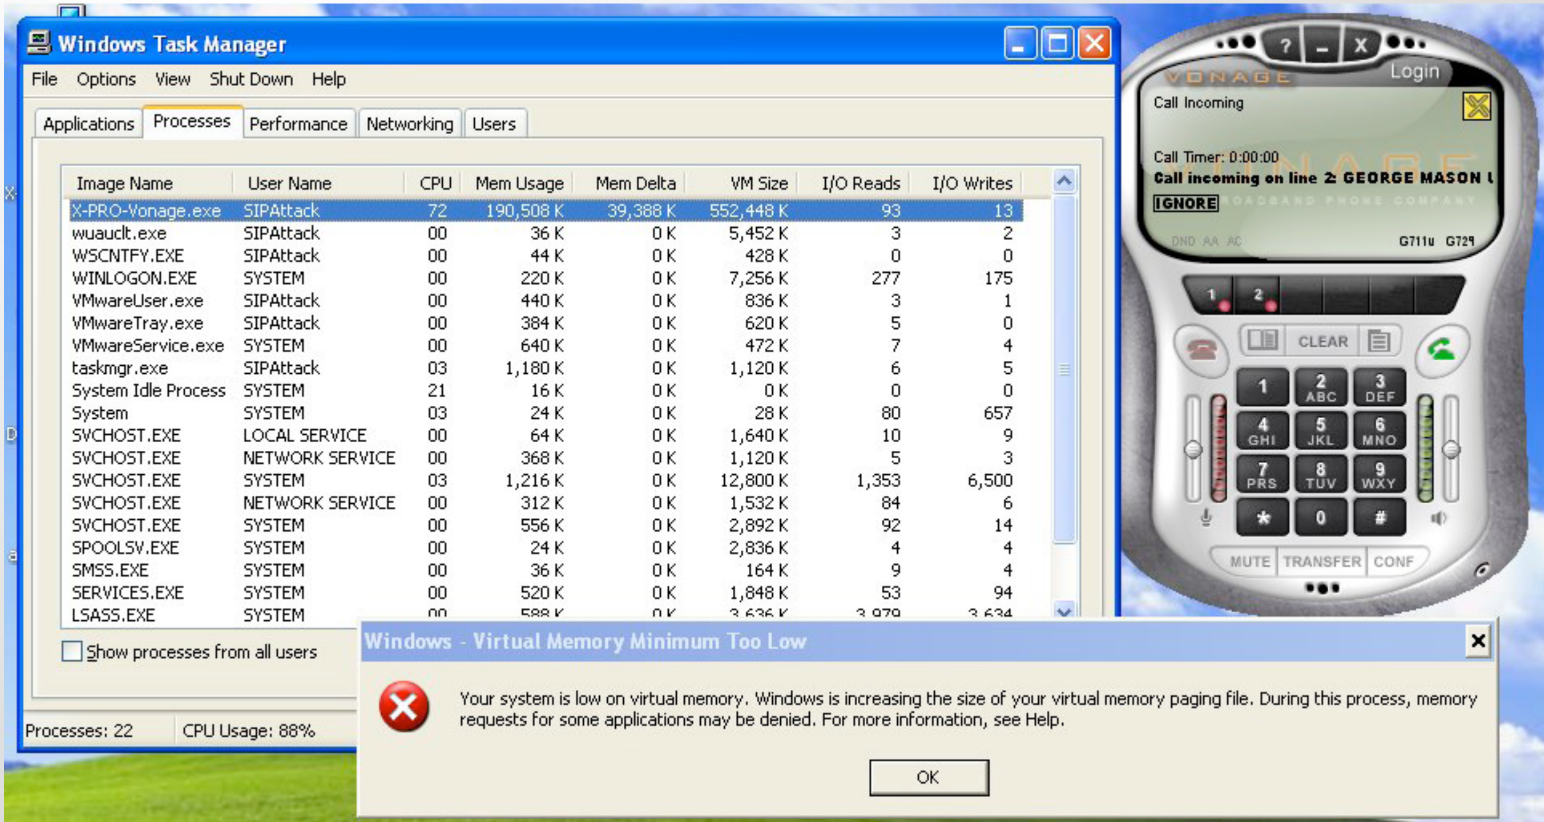

Noisy Attack

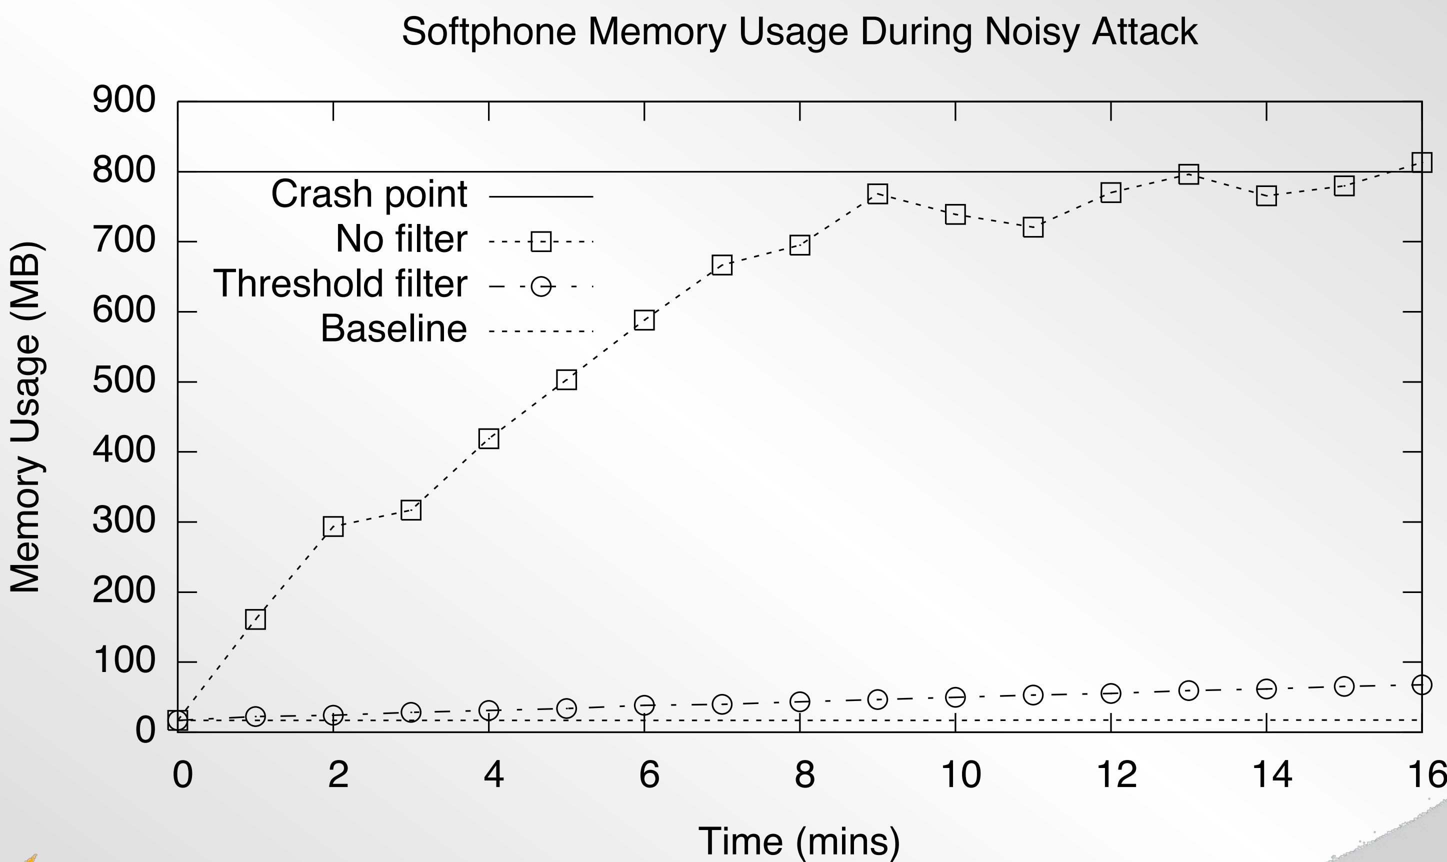

GEORGE MASON UNIVERSITY

Where Innovation Is Tradition

Stealthy Attack

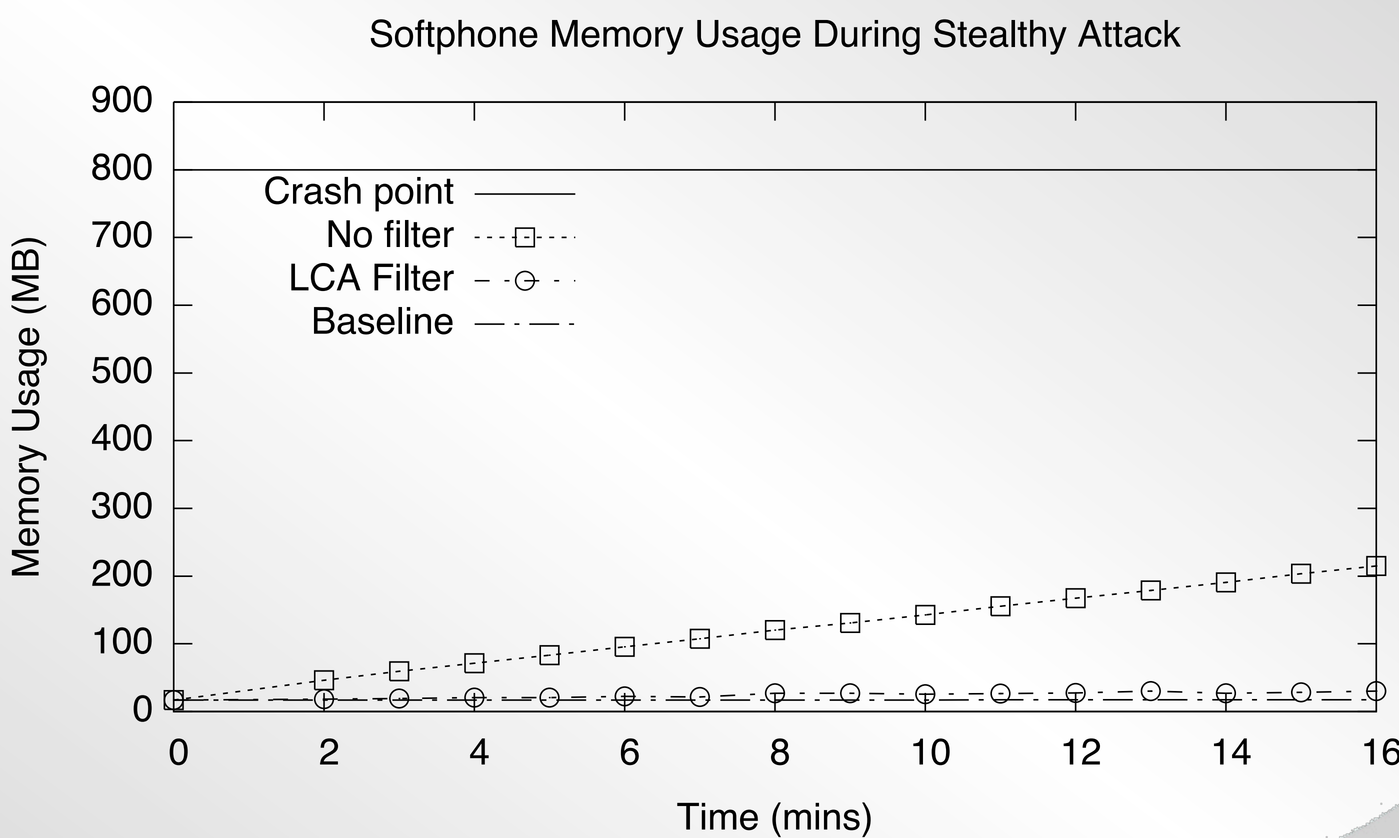

Defense Effectiveness

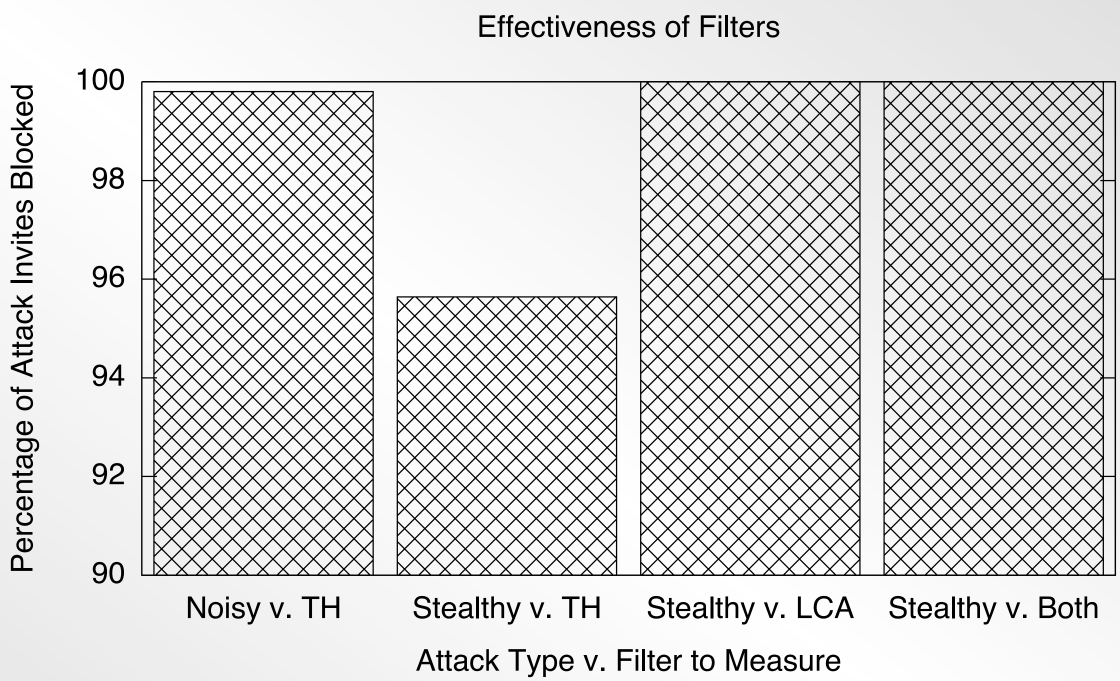

- Stealthy invites accounted for only 15.2% of packets against TH
- LCA tested with mixture of legitimate and illegitimate invites.
- ‘Both’ involves LCA feeding its output into TH

Defense Latency

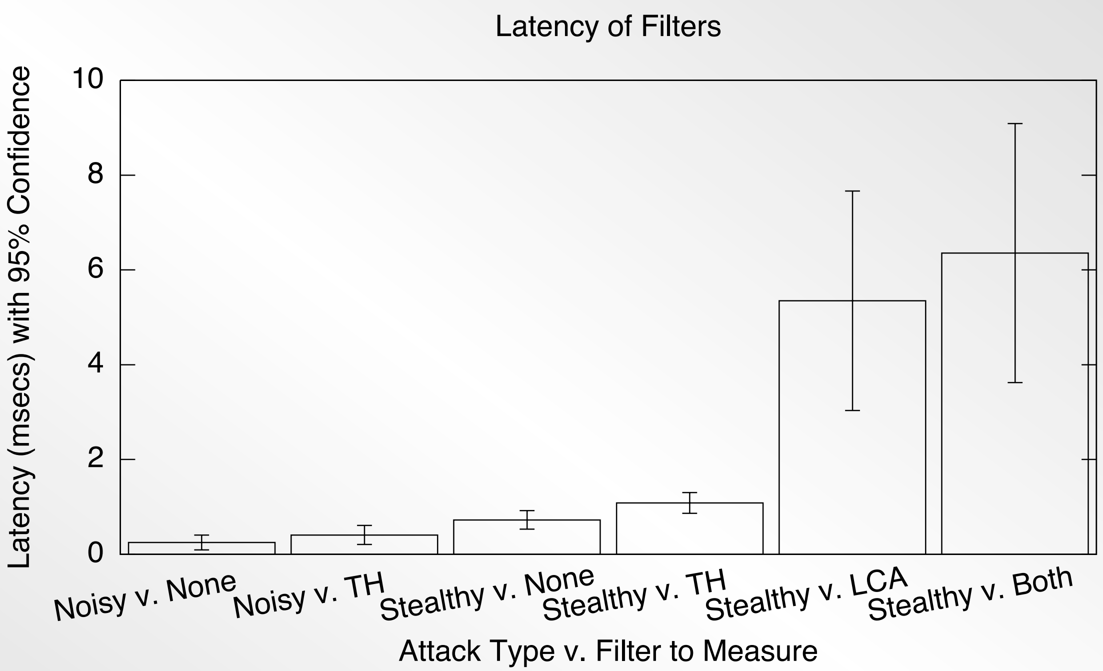

- Per RFC 2544
- TH introduces less than 1 millisecond, LCA less than 5 milliseconds
- No noticeable impact on VoIP signaling functionality observed

- Introduction
- Background
- Disabling the Softphone Host
- Defense
- Experiments
- **Conclusion**

Conclusion

- Features exploited are SIP, not Vonage
  - Enforcing SIP authentication could help mitigate
- First to demonstrate disabling the VoIP application host; via two attacks
  - Noisy attack effective in minutes
  - Stealthy attack only 1/(n+1) the noisy rate
- Presented packet filters to mitigate
  - Threshold: ultra-low overhead, highly effective
  - LCA: accurately drops stealthy attack from valid traffic

Thank you for your time

- Any questions?

Post conference, please contact Dr. Xinyuan Wang

- xwangc@gmu.edu

Invite Message

INVITE sip:17031234567@129.174.130.175:5060 SIP/2.0 Via: SIP/2.0/UDP 216.115.20.41:5061 Via: SIP/2.0/UDP 216.115.20.29:5060 Via: SIP/2.0/UDP 216.115.27.11:5060;branch=z9hG4bK8AE8A3914F0 From: "GMU" <sip: 17032345678@216.115.27.11>;tag=455412559 To: <sip: 17031234567@voncp.com> Call-ID: 58A8C0B-8D6F11DC-B8E18C7A-2083704C@216.115.27.11 CSeq: 101 INVITE Contact: <sip: 17032345678@216.115.20.41:5061> Max-Forwards: 13 X-Von-Relay: 216.115.27.30 Content-Type: application/sdp Content-Length: 361

v=0 o=CiscoSystemsSIP-GW-UserAgent 5330 7344 IN IP4 216.115.27.30 s=SIP Call c=IN IP4 216.115.27.30 t=0 0 m=audio 13598 RTP/AVP 0 18 2 100 101 c=IN IP4 216.115.27.30 a=rtpmap:0 PCMU/8000 a=rtpmap:18 G729/8000 a=fmtp:18 annexb=no a=rtpmap:2 G726-32/8000 a=rtpmap:100 X-NSE/8000 a=fmtp:100 192-194 a=rtpmap:101 telephone-event/8000 a=fmtp:101 0-16

Detailed Noisy Attack

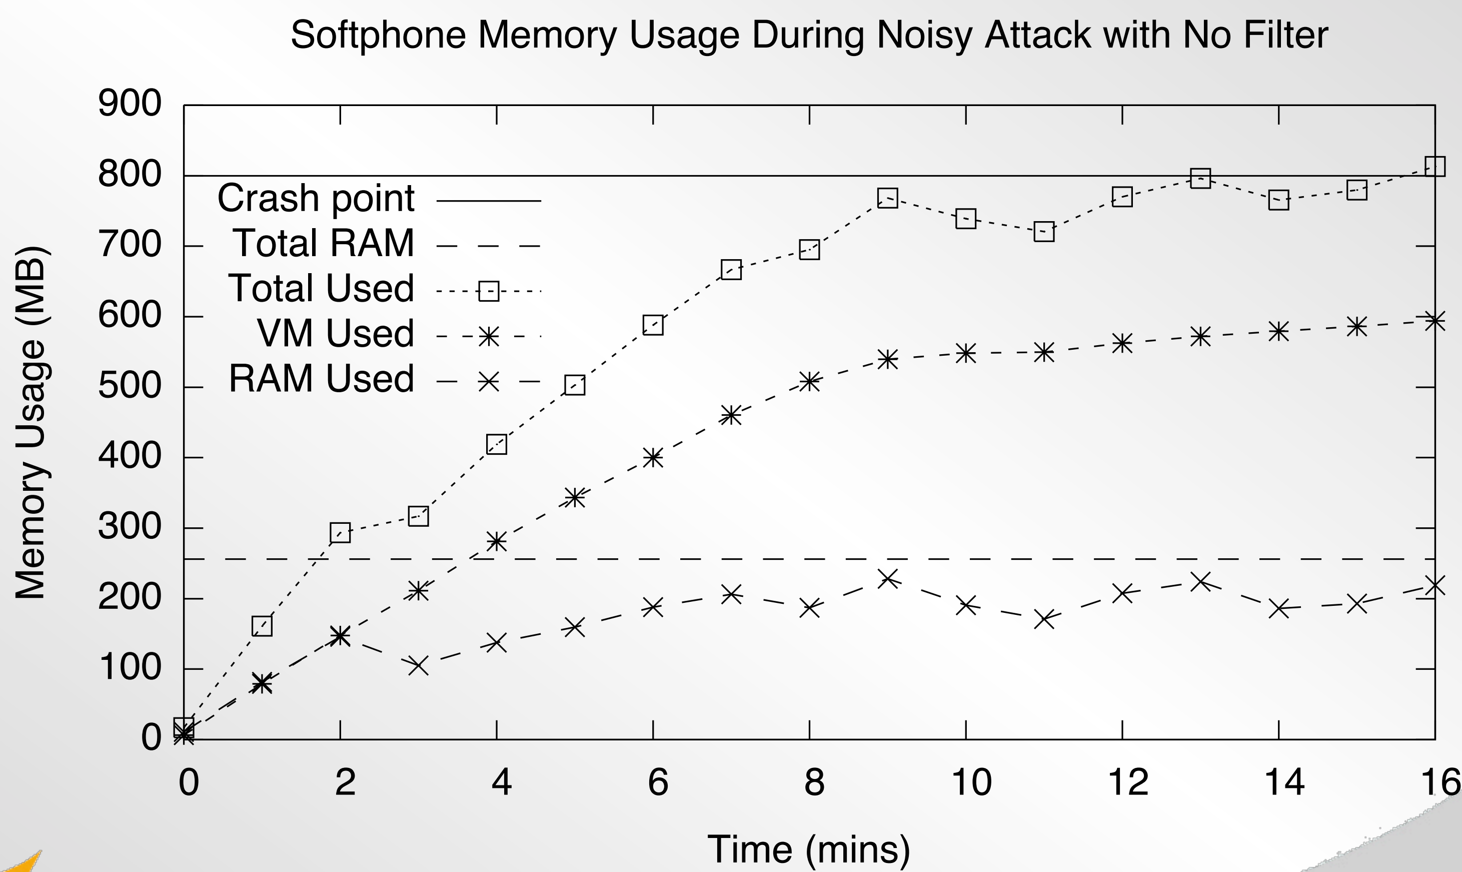

Defense Throughput

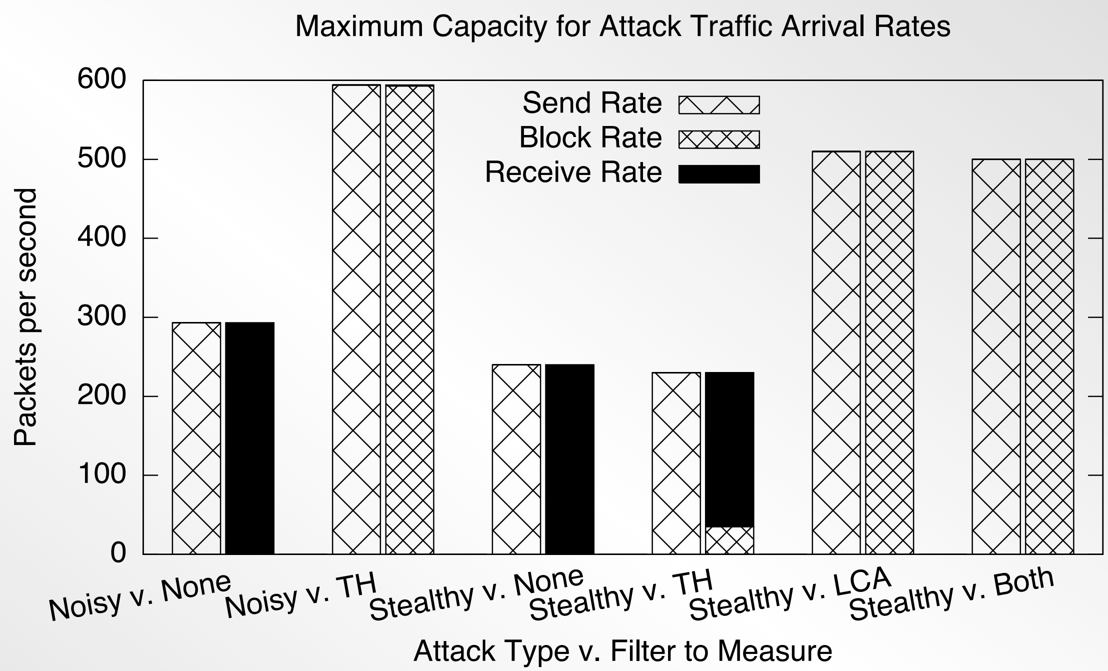

- Fastest packet rate without packet loss, RFC 2544
  - Slightly different since filtering drops packets (success if send = block + received)
  - Used to calculate latency