

# *Complexity Attacks on ipfw*

David Malone <dwmalone@freebsd.org>  
Josh Tobin <rjt@maths.tcd.ie>

20 September 2009

## Complexity Attacks

- Working on some code — too slow.
- You implement some clever algorithm — nice and fast.
- Algorithm Performance: *average* and *worst case*.
- E.g. Quicksort.
- Who controls what performance you get?

See “Denial of Service via Algorithmic Complexity Attacks” by Scott A Crosby and Dan S Wallach in 2003 Usenix Security Symposium.

## Hash Table

Lookup scheme to avoid cost of searching full list.

carrot zuchinni ...

...apricot apple

becomes:

*a* apricot, apple

*b* banana

*c* carrot

· ...

*z* zuchinni

Hash function  $h(x)$ , XOR typical. Cost:  $O(N) \rightarrow O(N/H)$ .

## *Algorithmic Attacks*

Suppose attacker controls keys.

*a* abduce, abducens, abducent, abduct,  
abduction, abductor ...

*b*

.

*z*

Attacker finds  $x_i$  so that  $h(x_1) = h(x_2) = \dots = h(x_i)$ .

## *Crosby and Wallach*

*Bro* XOR based hash to store state. Dual core Pentium drop traffic using 30s of 160kbps or 6 minutes of 16kbps.

*Perl* Hash used for (uh) Hashes. Insert 90k short strings. Usually took under 2s. With crafted input took almost 2 hours.

Also attacks on directory cache, Python, GLIB, Java, ...  
Related attacks on browsers, web servers, ...


## *Fixes?*

- Don't use hashes, choose good worst case performance.
- — can be hard to fit into existing code.
- Use cryptographic hashes.
- — output truncated
- Use keyed hash.
- — sometimes use misnomer “universal” .

## *ipfw Flow Lookup*

- Long standing firewall in FreeBSD.
- About 4.0 it grew stateful features.
- Lookup single flow by tuple (src IP, dst IP, src port, dst port).
- Uses hash table as optimisation.
- For IPv4 96 bits of input.
- For IPv6 288 bits of input.
- Note inexact flow matching different!


# Demonstration attack






<i>Xor</i>	$h \leftarrow 0$ foreach ( <i>byte</i> [ <i>i</i> ]) return <i>h</i>	$h \leftarrow h \oplus \text{byte}[i]$
<i>DJB2</i>	$h \leftarrow 5381$ foreach ( <i>byte</i> [ <i>i</i> ]) return <i>h</i>	$h \leftarrow 33 * h + \text{byte}[i]$
<i>XorSum</i>	$h \leftarrow 0$ foreach ( <i>byte</i> [ <i>i</i> ]) return <i>h</i>	$h \leftarrow h + (\text{byte}[i] \oplus K[i])$
<i>SumXor</i>	$h \leftarrow 0$ foreach ( <i>byte</i> [ <i>i</i> ]) return <i>h</i> ;	$h \leftarrow h \oplus (\text{byte}[i] + K[i])$
<i>Universal</i>	$h \leftarrow 0$ foreach ( <i>byte</i> [ <i>i</i> ]) return <i>h</i> mod 65537	$h \leftarrow h + K[i] * \text{byte}[i]$
<i>Pearson</i>	$h_1 \leftarrow h_2 \leftarrow 0$ foreach ( <i>byte</i> [ <i>i</i> ]) return $h_1 + h_2 * 256$	$h_1 \leftarrow T_1[\text{byte}[i] \oplus h_1]$ $h_2 \leftarrow T_2[\text{byte}[i] \oplus h_2]$
<i>MD5</i>	return two bytes of MD5(bytes)	
<i>SHA</i>	return two bytes of SHA(bytes)	

# Hash Chain Length



# CPU Cost



## *Other options*

Don't need to use hash.

*Tree* Use lexical order to insert into tree.

*Red/Black Tree* Tree balanced by colouring.

*Splay Tree* Moves frequently accessed to top.

*Treap* Tree balanced using random heap.


Tree is baseline (and subject to attack). Others are not (obviously) subject to attack.

## *Design Aims/Method*

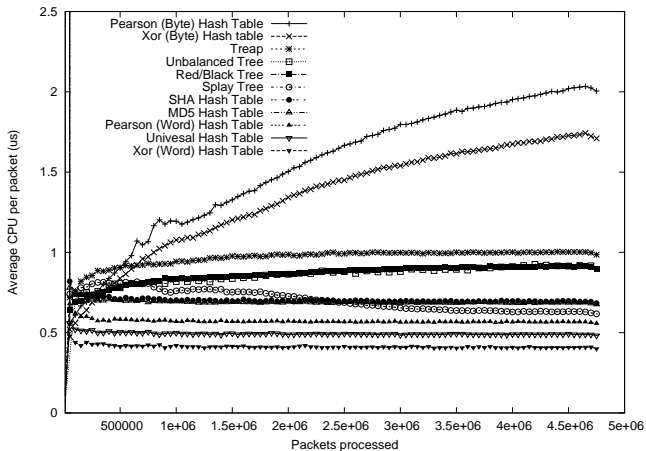
Want flow lookup to:

- Should perform OK under typical traffic.
- Should not degrade badly under attack.
- . . . typical performance depends on keys.
- . . . collect trace of traffic,
- . . . assess using pcap framework,
- . . . check performance in kernel.


# Traffic Trace



# Big CPU




# Small CPU





# Peak Forwarding



## *Conclusion*

- Take care choosing algorithms.
- XOR doesn't look so good.
- Universal hash actually pretty cheap.
- Looking at attacking some hashes:
- ... Simple Pearson.
- ... RSS Hash on Ethernet Cards.