

**Complexity Attack Resistant Flow Lookup  
Schemes for IPv6:  
A Measurement Based Comparison**

**David Malone and Josh Tobin.**

**2008-12-11**

# Hash Table

Lookup scheme to avoid cost of searching full list.

carrot zuchinni ...                      ... apricot apple

becomes:

**a** apricot, apple

**b** banana

• ...

**z** zuchinni

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

## Algorithmic Attacks

Worst case rather than typical behaviour. (Crosby and Wallach, 2003).

Suppose attacker controls keys.

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

**b**

.


**z**

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

## Flow Lookup

- Security applications often track IP flows.
- Lookup single flow by tuple (src IP, dst IP, src port, dst port).
- Hash table is one possible optimisation.
- For IPv4 96 bits of input.
- For IPv6 288 bits of input.
- Aim — `ipfw` for FreeBSD.
- Note inexact flow matching different!

# Demonstration attack



**Xor**     $h \leftarrow 0$

  foreach ( $byte[i]$ )

  return  $h$

$h \leftarrow h \oplus byte[i]$

**DJB2**     $h \leftarrow 5381$

  foreach ( $byte[i]$ )

  return  $h$

$h \leftarrow 33 * h + byte[i]$

**XorSum**     $h \leftarrow 0$

  foreach ( $byte[i]$ )

  return  $h$

$h \leftarrow h + (byte[i] \oplus K[i])$

**SumXor**     $h \leftarrow 0$

  foreach ( $byte[i]$ )

  return  $h$ ;

$h \leftarrow h \oplus (byte[i] + K[i])$

**Universal**     $h \leftarrow 0$

    foreach (*byte*[*i*])             $h \leftarrow h + K[i] * \textit{byte}[i]$

    return  $h \bmod 65537$

**Pearson**     $h_1 \leftarrow h_2 \leftarrow 0$

    foreach (*byte*[*i*])             $h_1 \leftarrow T_1[\textit{byte}[i] \oplus h_1]$


$h_2 \leftarrow T_2[\textit{byte}[i] \oplus h_2]$

    return  $h_1 + h_2 * 256$

**MD5**        return two bytes of MD5(bytes)


**SHA**        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

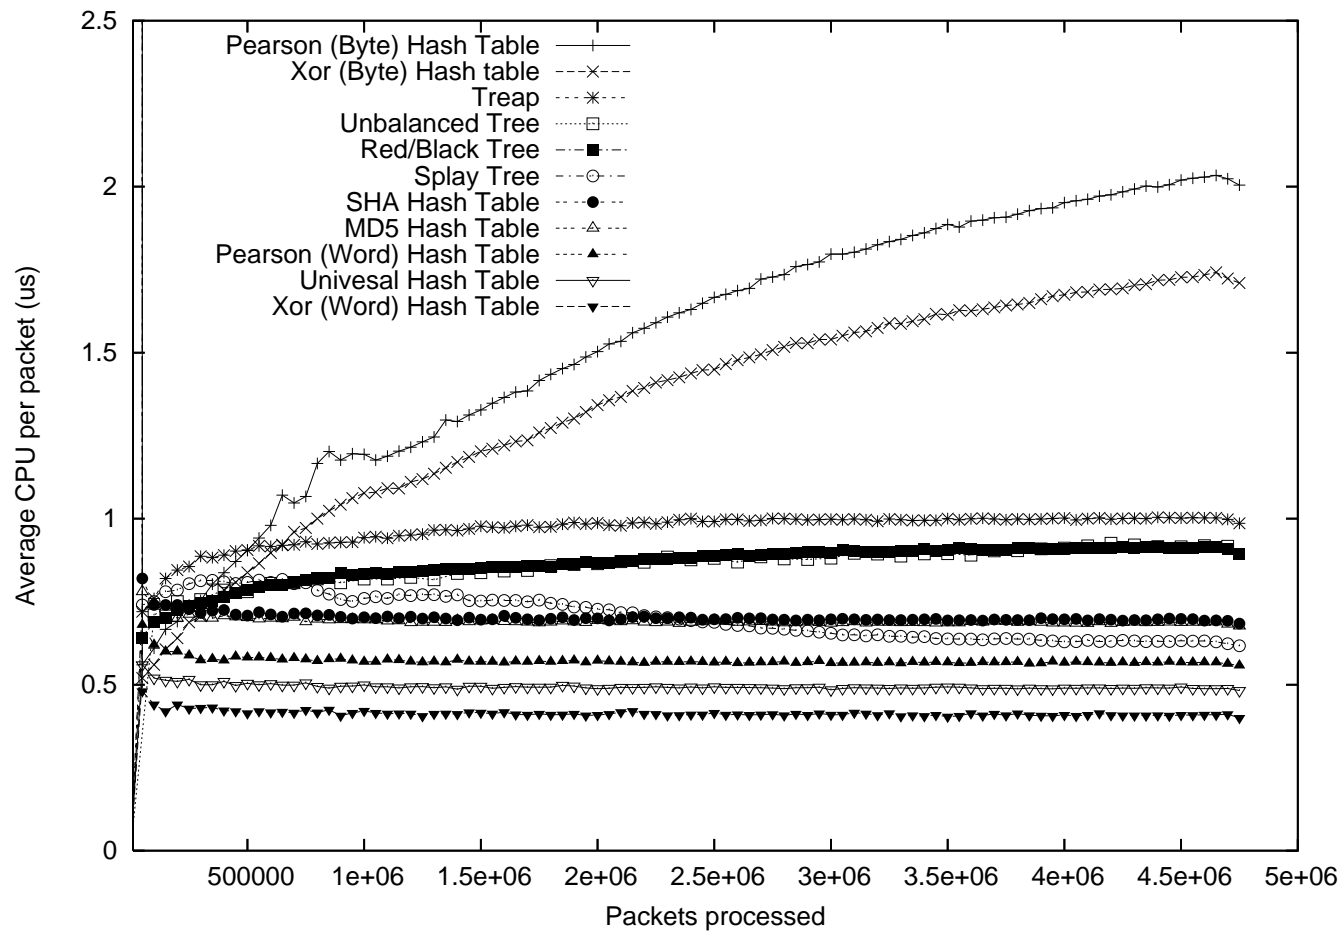
Want flow lookup to:

- 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



## Summary

- Looked at flow lookup schemes.
- Wanted attach resistant scheme,
- ... with good typical performance.
- Future: get code into FreeBSD.
- Future: look at attacks on hashes.
- Future: new hashing schemes.