

## 21:11 Windrose Fingerprinting of Code Architecture

by EVM

Often we come across a firmware from a device that we don't have in hand, and don't know anything about beyond pictures or sales glossies on a vendor website. We'd like to be able to load this firmware into a disassembler and analyze it anyway.

ELF firmware files will happily tell you and your disassembler the CPU architecture, but what do we do when analyzing a flat binary firmware file? We need a method to determine the architecture by comparing the file to previous samples from known architectures.

Each processor architecture has a unique byte histogram fingerprint, which others have described previously. This is because in machine code some types of opcodes are used more frequently than others (Register/memory move, comparison, jump/branch/call are usually the most common.) This gives each architecture a unique balance of bytes reflecting the designer's choice of representation of common and uncommon opcodes.

What I'm adding to the toolbox here is the concept of visualizing byte histograms as a windrose diagram. Byte histograms can be compared using a chi-squared test, but windrose diagrams may allow for a more-nuanced, visual comparison.

The following diagrams were generated from samples, mostly Linux kernels and Busybox binaries, and the occasional random large firmware file. Linux kernels and Busybox binaries work well because they are very large and contain a mix of lots of different kinds of code.

Here is a Python script that outputs a windrose diagram for a sample that you can compare against the fingerprints shown. This code bins the bytes in groups of four for more readable diagrams, and ignores bytes 0x00, 0x40, 0x80 and 0xC0 (to avoid over-representing top address bytes). Note that for best results you need to only map the text section of a binary, and remove any padding. Normally in a flat firmware binary the text section appears before the data section, and depending upon where you make the cut, your mileage may significantly vary on very small binaries.

As an example, Figure 12 presents windrose diagrams from the `.text` section of two 32-bit MIPS binaries. These are the first two MIPS binaries in the ALLSTAR dataset<sup>31</sup> whose `.text` section is greater than 64KB, `7kaa` from Debian's `7kaa` package, and `jmd1x` from Debian's `aajm` package. Notice their largest three spikes (the 0x00, 0x20, and 0x8C bins) match the 32-bit MIPS fingerprint well. The double spike (0x20 and 0x24 bins) appears in all three prints. `jmd1x` has a shorter spike at 0x00, and a longer spike at 0x08, but we can still easily see that its best match is 32-bit MIPS.

### THE HQ-140-X...



*"A Ham's Dream"*  
—says W4VPU

After trying out his new Hammarlund HQ-140-X receiver, Harry H. Harris, Jr., of Charlottesville, Va., W4VPU commented, "This is truly a Ham's dream."

Creating 'dream' equipment for hams is the Hammarlund goal. How well this goal has been achieved is proven by the enthusiastic comments received from satisfied Hams. They appreciate the little extras in design, circuitry and construction built into every Hammarlund product.

For example, the HQ-140-X—the amateur receiver built to professional standards—is rated XFB by Hams everywhere because of its—

**FREQUENCY STABILITY** — less than .01% frequency drift after warmup anywhere from 540 Kc. to 31 Mc.

**EXTREME SELECTIVITY** — sharp signal separation even in the most crowded bands.

**LOW NOISE LEVEL** — a noise limiter that really works.

**RUGGED CONSTRUCTION** — built for easy use for many years.

The HQ-140-X is available either as a cabinet model or for rack mounting. For complete details, write to The Hammarlund Manufacturing Co., Inc., 460 West 34th Street, New York 1, New York. Ask for Bulletin R-3.



<sup>31</sup><https://allstar.jhuapl.edu>




Figure 12: Fingerprints of two sample MIPS executables.

```

#!/usr/bin/python
2 import sys
  import struct
4
  fname = sys.argv[1]
6
  bytes=[]
8  entries=[]
  total=0
10
  for i in xrange(0,256):
12     bytes.append(0)
14
  with open(fname,'rb') as f:
    while True:
16         b=f.read(1)
        if b=="":
18             break
        bint = struct.unpack('<B',b)[0]
20         bytes[bint]+=1
22
  for i in xrange(0,256,4):
    entry=0
24     for j in xrange(0,4):
        if (((i+j) % 0x40) != 0):
26             entry+=int(bytes[i+j])
        entries.append(entry)
28     total+=entry
30
  for i in xrange(0,0x40):
    print "%f,%f" % (100*entries[i]/1.0/max(entries), i*360/1.0/0x40)

```

