Cisco Systems

PRNG in IOS

Gaus – PSIRT IM

<gaus@cisco.com>

Overview

Cisco.com

- How it started
- What it looked like
- How it was improved
- How was tested
- Possible further improvements

[redacted]

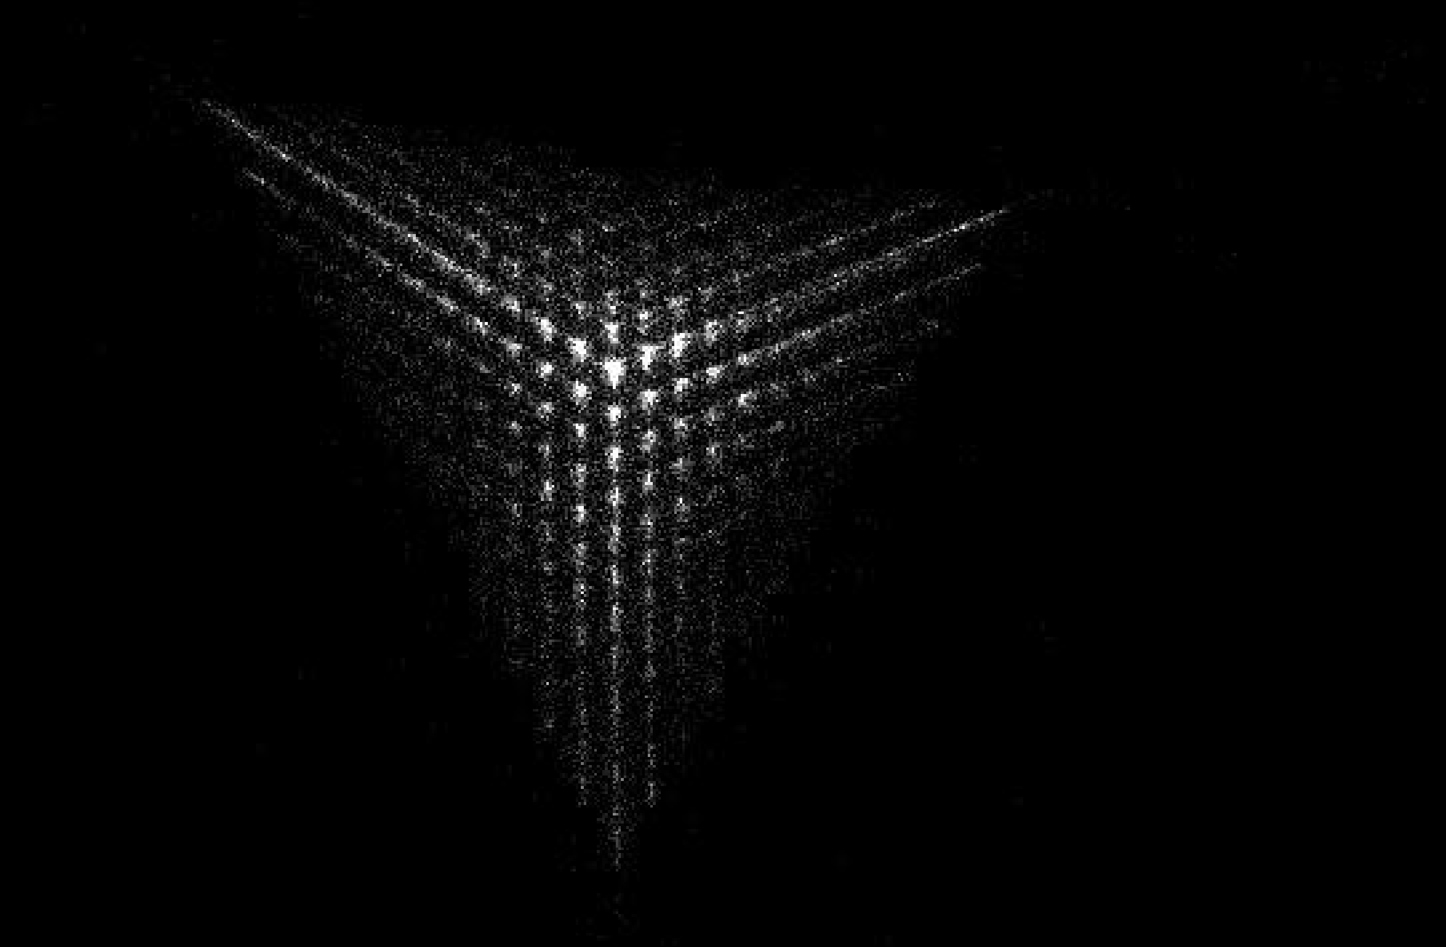

What were looking at

- “Big Bang”, who knows how it looked like?
- Starry night, Copyright NASA, STScI, HubbleSite
- “Strange Attractors and TCP/IP Sequence Number Analysis” by Michal Zalewski

Beginning

- **It started by ISNs in TCP session**
- **They should be unpredictable but they were not**

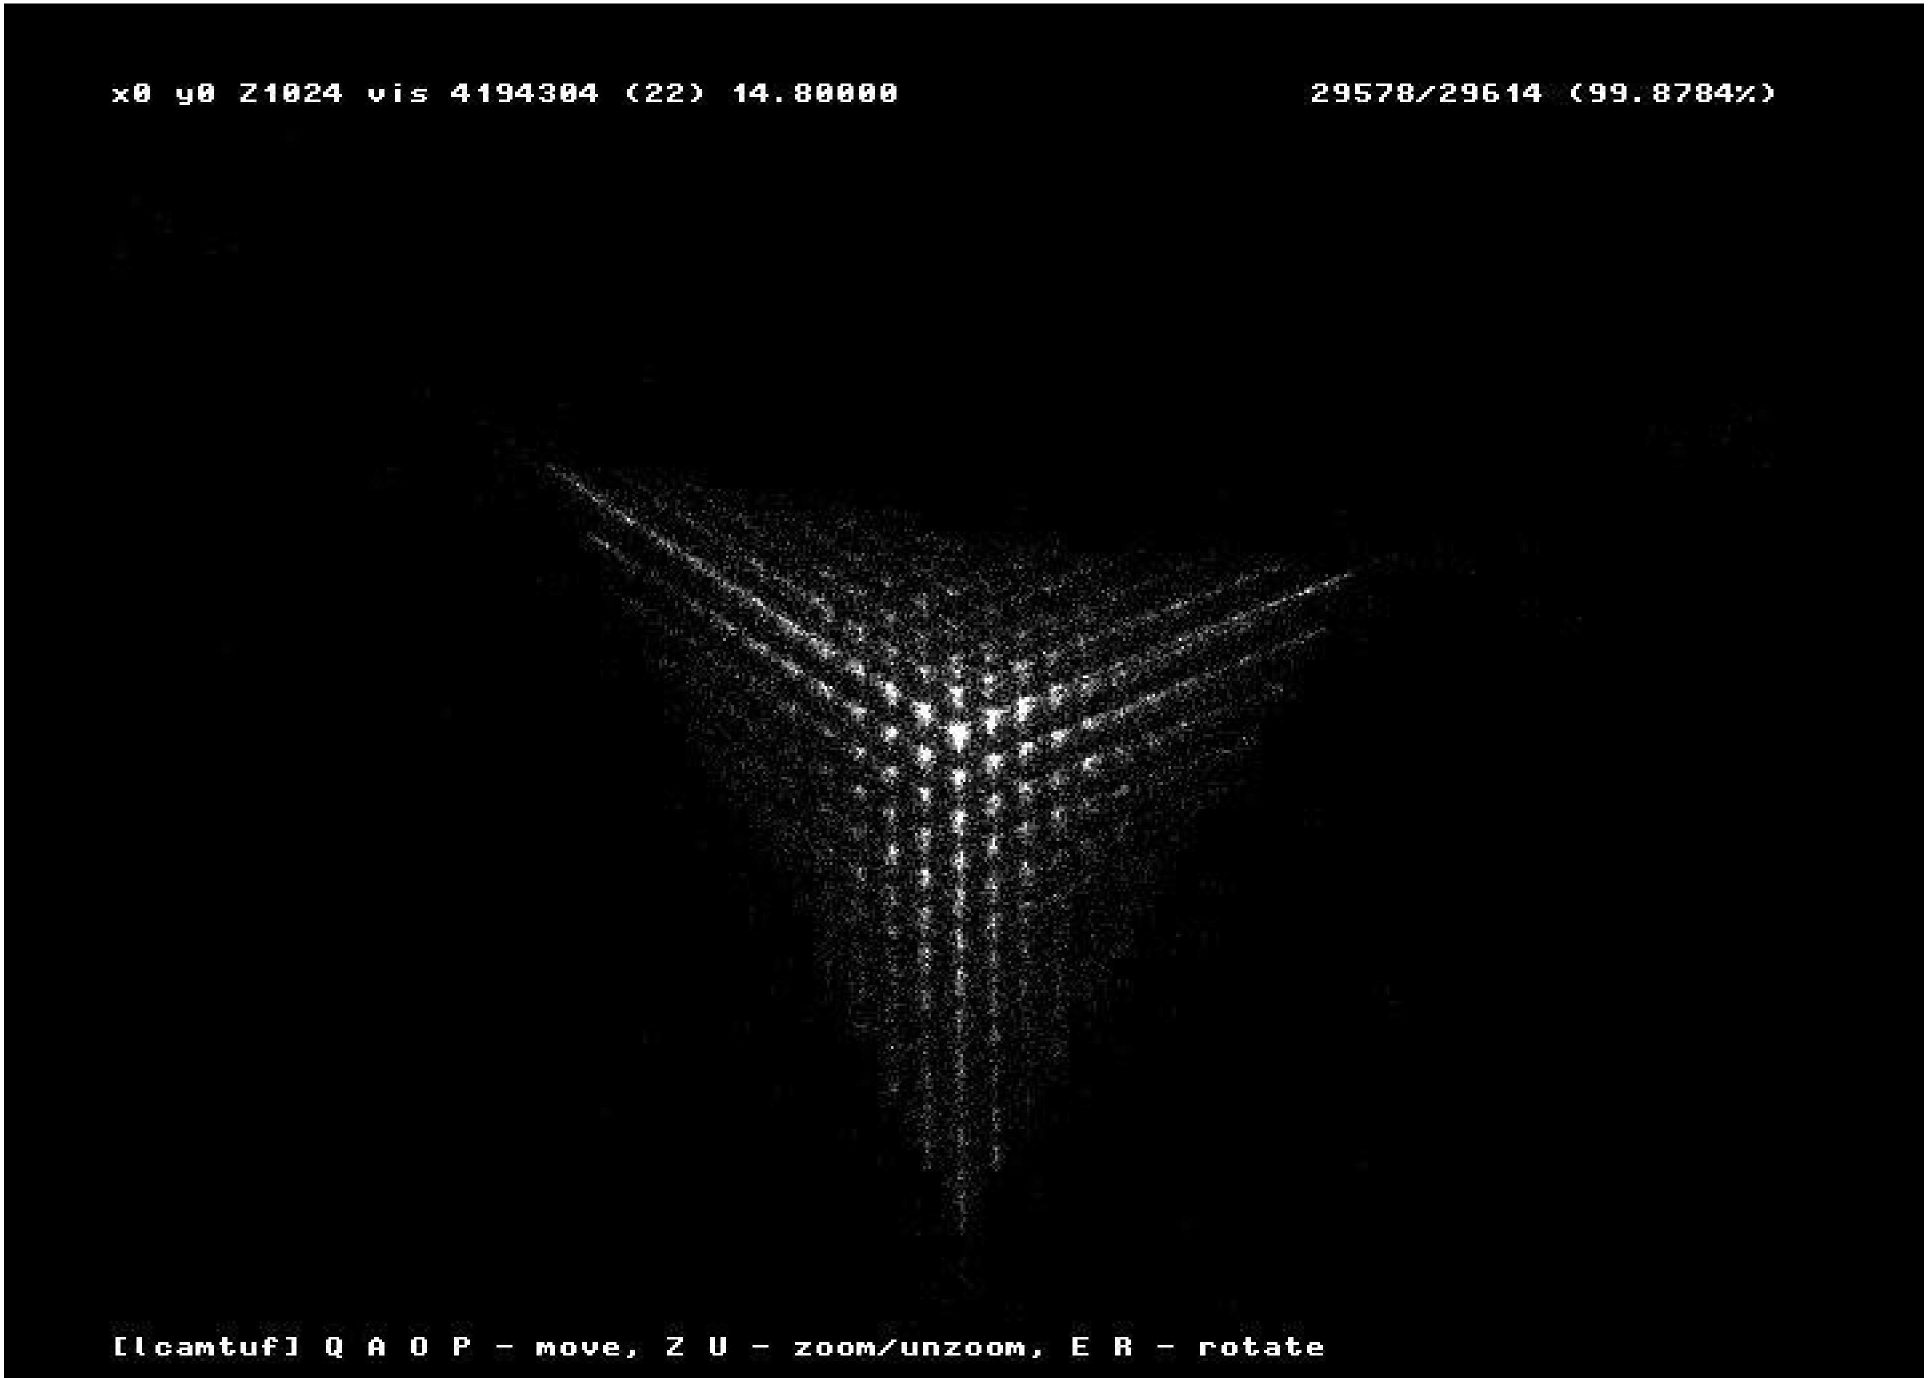
x0 y0 Z1024 vis 4194304 (22) 14.80000
29578/29614 (99.8784%)
[lcamtuf] Q A O P - move, Z U - zoom/unzoom, E R - rotate

Improving ISNs

- One way is to make ISN as random as possible
- What exactly "random" means?
- Unpredictable and next-bit test
- The existing PRNG was not adequate for the purpose
- The solution is to introduce a new one

How PRNG works

- The universal recipe is the same:
  - ➢Take some fresh entropy and put it into a pot
  - ➢Add more entropy whenever you have a chance and stir it in
  - ➢Serve when needed but not forget to stir

Some mixing tools

SHA

MD5

AES

DES

More mixing tools

The challenge

- **Where to find entropy**
- **IOS is closed system and it does not have:**
  - **Hard disk**
  - **Mouse**
  - **Keyboard**

Some unsuccessful ideas

- MAC or IP addresses
- Packet length, timing between packets
- Environmental temperature
- CPU fan rotation
- Wireless noise, microphone, camera
- “Something” from the memory

More promising ideas

- truerandom() function but IOS is not preemptive
- Timing between consecutive passes in a simple loop
- Time when the function is invoked
- “something” from the memory

How it looks today

- PRNG uses GF-based mixing function and it is extracted using MD5
- Entropy is slow to accumulate
- PRNG passes all statistical tests


x0 y0 Z1 vis 2147483648 (32) 0.00000
64389/100000 (64.3890%)
[lcamtuf] Q A O P - move, Z U - zoom/unzoom, E R - rotate

How to test a sequence

- **How random is your “random” sequence?**
- **Is “1111111111111111111111” more random than “0101010101010101010101” or “10011010100101101”?**
- **We can only test for statistical properties of a sequence.**

Tools used for testing

- Diehard
- NIST Statistical Test Suite
- Some others were tried but were not adequate

Diehard

- **Not really user friendly**
- **Need some knowledge to interpret the results**
- **Very powerful**
- **Needs large input (~8*10^9 bits)**

NIST STS

- Nicer interface
- Sometimes can be hard to select right parameters and input sequence length
- An par with Diehard

A sample of Diehard output

Cisco.com

```
::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
::            This is the BIRTHDAY SPACINGS TEST                  ::
:: Choose m birthdays in a year of n days.  List the spacings     ::
:: between the birthdays.  If j is the number of values that      ::
:: occur more than once in that list, then j is asymptotically    ::
:: Poisson distributed with mean m^3/(4n).  Experience shows n    ::
:: must be quite large, say n>=2^18, for comparing the results    ::
:: to the Poisson distribution with that mean.  This test uses    ::
:: n=2^24 and m=2^9,  so that the underlying distribution for j   ::
:: is taken to be Poisson with lambda=2^27/(2^26)=2.  A sample    ::
:: of 500 j's is taken, and a chi-square goodness of fit test     ::
:: provides a p value.  The first test uses bits 1-24 (counting   ::
:: from the left) from integers in the specified file.            ::
::   Then the file is closed and reopened. Next, bits 2-25 are    ::
:: used to provide birthdays, then 3-26 and so on to bits 9-32.   ::
:: Each set of bits provides a p-value, and the nine p-values     ::
:: provide a sample for a KSTEST.                                 ::
::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
```

A sample of Diehard output (cont.)

Cisco.com

```
BIRTHDAY SPACINGS TEST, M= 512 N=2**24 LAMBDA=  2.0000
           Results for newrnd.bin
                   For a sample of size 500:     mean
           newrnd.bin      using bits  1 to 24   2.000
 duplicate       number       number
 spacings       observed     expected
        0          65.       67.668
        1         137.      135.335
        2         138.      135.335
        3          93.       90.224
        4          41.       45.112
        5          15.       18.045
  6 to INF         11.        8.282
 Chisquare with  6 d.o.f. =     2.04 p-value= 0.084410
```

A sample of STS output

Cisco.com

Statistical Test	P-value
Frequency	0.604458
Block Frequency (m = *100*)	0.833026
Cusum-Forward	0.451231
Cusum-Reverse	0.550134
Runs	0.309757
Long Runs of Ones (M = *10000*)	0.657812
Rank	0.577829
Spectral DFT	0.086702
NonOverlapping Templates (m = *9*, B = *000000001*)	0.496601
Overlapping Templates (m = *9*)	0.339426
Universal (L = *7*, Q = *1280*)	0.411079
Approximate Entropy (m = *5*)	0.731449
Random Excursions (x = *+1*)	0.000000
Random Excursions Variant (x = *-1*)	0.000000
Lempel Ziv Complexity	0.398475
Linear Complexity (M = *500*)	0.309412
Serial (m = *5*, $\nabla\Psi^2_m$)	0.742275

Possible improvements

- The current PRNG is not the fastest in the block
- Possible replacements with AES-based
- Retaining entropy over reloads

Links

Cisco.com

- http://razor.bindview.com/publish/papers/tcpseq.html
- http://lcamtuf.coredump.cx/newtcp/
- http://www.cs.berkeley.edu/~daw/rnd/mab-rand
- http://www.schneier.com/yarrow.html

More links

- http://csrc.ncsl.nist.gov/rng/
- http://stat.fsu.edu/pub/diehard/

Cisco Systems