

# Mobile Measurement of Path Transparency

Iain R. Learmonth   Gorry Fairhurst  
{iain,gorry}@erg.abdn.ac.uk

University of Aberdeen  
MAMI Project

2nd Workshop on Mobile Network Measurement (MNM'18)  
June 25, 2018  
Vienna, Austria

*This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 688421. The opinions expressed and arguments employed reflect only the authors' view. The European Commission is not responsible for any use that may be made of that information.*

# Active Measurement of Path Transparency

- Methodology:
  - 1 Throw packets at the Internet
  - 2 See what happens
- Ideal: two-ended A/B testing
- Scalable: one-ended A/B testing
- Multiple sources: **isolate** on-path from near-target impairment




Figure: One-Ended vs. Two-Ended Testing

# History

*ecns spider*

- The original implementation supported by mPlane/RITE
- Three distinct components:
  - DNS List Resolver
  - QoF Flow Meter
  - Active Traffic Generator
- Used hardcoded `sysctl(1)` and `iptables(1)` commands to cause packets to be emitted with various ECN-related flags
- Source code:  
<https://github.com/britram/ecns spider>




Figure: Original Architecture

# History

## *ecns spider* Results

- ECN negotiation was found to be successful for 56.17% of hosts connecting for IPv4, 65.41% for IPv6, from the Alexa top 1 million list [4]
- This continues a trend ETH started observing with *ecns spider* in 2013 [2]




Figure: ECN Support in the Alexa Top 1 Million

# PATHspider 1.0

<https://github.com/mami-project/pathspider/tree/1.0.1>

- Architecture based closely on the original ecnspider
- Generalised to support more than just ECN
  - Added TCP Fast Open and DiffServ Codepoints
- Still performing A/B testing, but with more A/B tests
- Replaced QoF with a Python flowmeter implementation using python-libtrace
- Began to develop a generalised measurement methodology for path transparency testing
- Published at 2016 Applied Networking Research Workshop [3]

# Plugin Architecture

- The plugin architecture was not as generalised as it could have been
- Plugin methods:
  - `config_zero`
  - `config_one`
  - `connect`
  - `post_connect`
  - `create_observer`
  - `merge`

# Built-In Flowmeter

- PATHspider's built in flow meter is extensible via the plugin architecture
- Using python-libtrace to dissect packets, any flow property imaginable can be reported back based on the raw packets:
  - ECN negotiation (IP/TCP headers)
  - Bleaching of bits, dropping of options
  - Checksum recalculations

# PATHspider 1.0 Results

We presented some initial findings along with the publication of PATHspider 1.0 [3]:


## Explicit Congestion Notification (ECN)

State of ECN server-side deployment, as measured from a

Digital Ocean vantage point in Amsterdam on 13th June 2016:

	IPv4	IPv6	all
<b>No ECN connectivity issues</b>	99.5%	99.9%	99.5%
<b>ECN successfully negotiated</b>	70.0%	82.8%	70.5%

ECN negotiation by Alexa rank bin:



## DiffServ Code Points (DSCP)

Initial study: 10,006 of 96,978 (10.31%) of Alexa Top 100k

websites had unexpected, non-zero DSCP values. More

measurement was needed to better characterize these anomalies.

## TCP Fast Open (TFO)

Initial study: 330 IPv4 and 32 IPv6 addresses of Alexa Top 1M

are TFO-capable (of which 278 and 28 respectively are Google

properties). DDoS prevention services, enterprise firewalls, and

CPE tend to interfere with TFO. More measurement was

necessary to analyze impairments.



# PATHspider 2.0

- Architecture changed to add a flow combiner
- Generalised to support more than just A/B testing
  - Any permutation of any number of tests
- Replaced PATHspider's HTTP code with cURL
- Added framework for packet forging based plugins using Scapy
- Completely rewritten (in Go) target list resolver
- Observer modules usable for standalone passive observation or analysis
- Source code: <https://github.com/mami-project/pathspider/tree/2.0.0/>

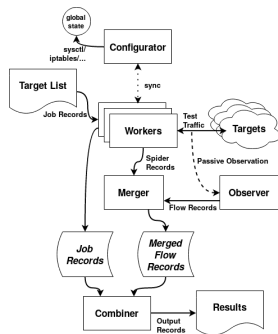


Figure: New Architecture

# Plugin Types

- Synchronised (traditional ecnspider)
  - ECN, DSCP
- Desynchronised (traditional ecnspider, no configurator)
  - TFO, H2, TLS NPN/ALPN
- Forge (new in PATHspider 2.0!)
  - Evil Bit, UDP Zero Checksum, UDP Options
- Single (new, and fast)
  - Various TCP Options

# Connection Helpers

- Instead of writing client code, use the code that already exists
- In the `pathspider.helpers` module:
  - DNS (`dnslib`)
  - HTTP/HTTPS (`pycURL`)
  - TCP (Python socket)
- For synchronised plugins, just use the helper
- For desynchronised plugins, the helpers are customisable, e.g. `cURL` helpers accept arbitrary `CURLOPTs`

# Synchronized Plugin

- SynchronizedSpider plugins use **built-in connection methods** along with **global system configuration** to change the behaviour of the connections
- Configuration functions are at the heart of a SynchronizedSpider plugin
- Configuration functions may make calls to sysctl or iptables to make changes to the way that traffic is generated.
- One function should be written for each of the configurations and PATHspider will ensure that the configurations are set before the corresponding traffic is generated. It is the responsibility of plugin authors to ensure that any configuration is reset by the next configuration function if that is required

# Synchronized Plugin

```
1 class ECN(SynchronizedSpider, PluggableSpider):
2     def config_no_ecn(self): # pylint: disable=no-self-use
3         """
4         Disables ECN negotiation via sysctl.
5         """
6
7         logger = logging.getLogger('ecn')
8         subprocess.check_call(
9             ['/sbin/sysctl', '-w', 'net.ipv4.tcp_ecn=2'],
10            stdout=subprocess.DEVNULL,
11            stderr=subprocess.DEVNULL)
12        logger.debug("Configurator disabled ECN")
13
14    def config_ecn(self): # pylint: disable=no-self-use
15        """
16        Enables ECN negotiation via sysctl.
17        """
18
19        logger = logging.getLogger('ecn')
20        subprocess.check_call(
21            ['/sbin/sysctl', '-w', 'net.ipv4.tcp_ecn=1'],
22            stdout=subprocess.DEVNULL,
23            stderr=subprocess.DEVNULL)
24        logger.debug("Configurator enabled ECN")
25
26    configurations = [config_no_ecn, config_ecn]
```

Listing 1: Configuration Functions for the ECN Plugin

# Desynchronized Plugin

- DesynchronizedSpider plugins modify the connection logic in order to change the behaviour of the connections. There is no global state synchronisation and so a DesynchronizedSpider can be more efficient than a SynchronizedSpider
- Connection functions are at the heart of a DesynchronizedSpider plugin
- These use a connection helper (or custom connection logic) to generate traffic towards with a target to get a reply from the target
- One function should be written for each connection to be made, usually with at least two functions to provide a baseline followed by an experimental connection

# Desynchronized Plugin

```
class H2(DesynchronizedSpider, PluggableSpider):
2   def conn_no_h2(self, job, config): # pylint: disable=unused-argument
        curlopts = {}
4       curlinfos = [pycurl.INFO_HTTP_VERSION]
        if self.args.connect == "http":
6           return connect_http(self.source, job, self.args.timeout, curlopts, curlinfos)
        if self.args.connect == "https":
8           return connect_https(self.source, job, self.args.timeout, curlopts, curlinfos)
    )
    else:
        raise RuntimeError("Unknown connection mode specified")

12  def conn_h2(self, job, config): # pylint: disable=unused-argument
        curlopts = {pycurl.HTTP_VERSION: pycurl.CURL_HTTP_VERSION_2_0}
14       curlinfos = [pycurl.INFO_HTTP_VERSION]
        if self.args.connect == "http":
16           return connect_http(self.source, job, self.args.timeout, curlopts, curlinfos)
        if self.args.connect == "https":
18           return connect_https(self.source, job, self.args.timeout, curlopts, curlinfos)
    )
    else:
20       raise RuntimeError("Unknown connection mode specified")

22  connections = [conn_no_h2, conn_h2]
```

Listing 2: Connection Functions for the H2 Plugin

# Forge Plugin

- ForgeSpider plugins use Scapy to send forged packets to targets
- The heart of a ForgeSpider is the `forge()` function
- This function takes two arguments, the job containing the target information and the sequence number
- This function will be called the number of times set in the packets metadata variable and `seq` will be set to the number of times the function has been called for this job



# Single Plugin

- SingleSpider uses the built-in connection helpers to make a single connection to the target which is optionally observed by Observer chains
- This is the simplest model and only requires a `combine_flows()` function to generate conditions from the connection helper output and flow record output from the Observer

# Observer Modules

- While these used to be part of plugins in PATHspider 1.0, they are now independent and so can be reused across multiple plugins:
  - BasicChain, DNSChain, DSCPChain, ECNChain, EvilChain, ICMPChain, TCPChain, TFOChain
- These can also be used together, limiting each chain to just a single layer and letting the combiner produce conditions
- Chains can produce information to be consumed by other chains later in the list
- These can be used independently of a PATHspider measurement:

```
irl@z~$ pspdr observe tcp ecn
```

Listing 3: Running the PATHspider Observer independently

# Target List Resolution

- Hellfire is a parallelised DNS resolver. It is written in Go and for the purpose of generating input lists to PATHspider, though may be useful for other applications
- Can use many sources for inputs:
  - Alexa Top 1 Million Global Sites
  - Cisco Umbrella 1 Million
  - Citizen Lab Test Lists
  - OpenDNS Public Domain Lists
  - Comma-Separated Values Files
  - Plain Text Domain Lists



# Target List Resolution

Using *hellfire*

```
1 irl@z:~$ hellfire
Usage:
3  hellfire --topsites [--file=<filename>] [--output=<individual|array|oneeach>] [--type=<
   host|ns|mx>] [--canid=<canid address>]
   hellfire --cisco [--file=<filename>] [--output=<individual|array|oneeach>] [--type=<
   host|ns|mx>] [--canid=<canid address>]
5  hellfire --citizenlab [--country=<cc>|--file=<filename>] [--output=<individual|array|
   oneeach>] [--type=<host|ns|mx>] [--canid=<canid address>]
   hellfire --opendns [--list=<name>|--file=<filename>] [--output=<individual|array|
   oneeach>] [--type=<host|ns|mx>] [--canid=<canid address>]
7  hellfire --csv --file=<filename> [--output=<individual|array|oneeach>] [--type=<host|ns
   |mx>] [--canid=<canid address>]
   hellfire --txt --file=<filename> [--output=<individual|array|oneeach>] [--type=<host|ns
   |mx>] [--canid=<canid address>]
```

Listing 4: hellfire's Usage Help

```
irl@z~$ hellfire --cisco
```

Listing 5: Start Resolving the Cisco Umbrella List

# Packet Forging

- PATHspider uses the Scapy library for Python for packet forging
- This is the most flexible method of creating new measurement plugins for PATHspider

# Make a Packet

- Scapy packets are constructed layer by layer
- While you can specify raw bytes, Scapy provides a number of useful classes for common protocols, which makes things a lot easier

# Make a Packet

- Scapy must be launched with `sudo` as we will need to use "raw" sockets to emit forged packets.

```
1 irl@z:~$ sudo scapy3
3
5      aSPY//YASa
6      apyyyyCY//////////YCa
7      sY////////YSpcs  scpCY//Pp
8  ayp ayyyyyySCP//Pp      syY//C
9  AYAsAYYYYYYYY//Ps      cY//S
10     pCCCCY//p      cSSps y//Y
11     SPPPP//a      pP///AC//Y
12     A//A      cyP////C
13     p///Ac      sC///a
14     P////YCpc      A//A
15     scccccP///pSP///p      p//Y
16     sY////////y caa      S//P
17     cayCyayP//Ya      pY/Ya
18     sY/PsY////YCc      aC//Yp
19     sc  sccaCY//PCyapaapyCP//YSs
20     spCPY////////YPSps
21     ccaacs
>>>
|
| Welcome to Scapy
| Version 2.4.0
|
| https://github.com/secdev/scapy
|
| Have fun!
|
| We are in France, we say Skappee.
| OK? Merci.
|
| — Sebastien Chabal
|
using IPython 5.5.0
```

Listing 6: Launching Scapy

# Make a Packet

## IPv4 Header - Create and Dissect

```
1 >>> IP()
2 <IP |>
3 >>> i = IP()
4 >>> i.summary()
5 '127.0.0.1 > 127.0.0.1 hopopt'
6 >>> i.display()
7 ###[ IP ]###
8 version= 4
9 ihl= None
10 tos= 0x0
11 len= None
12 id= 1
13 flags=
14 frag= 0
15 ttl= 64
16 proto= hopopt
17 chksum= None
18 src= 127.0.0.1
19 dst= 127.0.0.1
   \options\
```

Listing 7: Creating and Dissecting an IPv4 Header



# Make a Packet

## IPv4 Header - Customize

```
>>> i = IP(src="192.0.2.1",dst="198.51.100.1",ttl=10)
2 >>> i
<IP ttl=10 src=192.0.2.1 dst=198.51.100.1 |>
4 >>> i.summary()
'192.0.2.1 > 198.51.100.1 hopopt'
6 >>> i.display()
###[ IP ]###
8   version= 4
   ihl= None
10  tos= 0x0
   len= None
12  id= 1
   flags=
14  frag= 0
   ttl= 10
16  proto= hopopt
   chksum= None
18  src= 192.0.2.1
   dst= 198.51.100.1
20  \options\
```

Listing 8: Customizing an IPv4 Header

# Make a Packet

## IPv4 Header - Bonus: Export a Dissection

```
>>> i.pdfdump()
```

### Listing 9: Create a PDF Export of a Dissection of the IP Header




Figure: PDF Export of IP Header Dissection

# Make a Packet

## TCP Header: Create and Dissect

```
1 >>> TCP()
  <TCP |>
3 >>> t = TCP()
  >>> t.summary()
5 'TCP ftp_data > http S'
  >>> t.display()
7 ###[ TCP ]###
  sport= ftp_data
  dport= http
  seq= 0
  ack= 0
  dataofs= None
  reserved= 0
  flags= S
  window= 8192
  chksum= None
  urgptr= 0
  options= []
```

Listing 10: Creating and Dissecting a TCP Header

# Make a Packet

## TCP Header: Customizing

```
1 >>> t = TCP(dport=443)
2 >>> t
3 <TCP dport=https |>
4 >>> t.summary()
5 'TCP ftp_data > https S'
6 >>> t.display()
7 ###[ TCP ]###
8     sport= ftp_data
9     dport= https
10    seq= 0
11    ack= 0
12    dataofs= None
13    reserved= 0
14    flags= S
15    window= 8192
16    chksum= None
17    urgptr= 0
    options= []
```

Listing 11: Customizing a TCP Header

# Make a Packet

## Sticking the Pieces Together

- The / operator is used to join layers together.
- Scapy will automatically set fields, such as the IP Protocol field, when you do this.
- When dissecting, Scapy will automatically choose the dissector to use based on fields such as the IP Protocol field.

```
>>> p=i/t
2 >>> p.summary()
  'IP / TCP 192.0.2.1:ftp_data > 198.51.100.1:https S'
4 >>> p.display()
  [... output snipped ...]
```

Listing 12: Sticking the IP and TCP Headers Together

# Make a Packet

View in Wireshark

```
1 >>> wrpcap("/tmp/scapy.pcap", [p])
```

Listing 13: Exporting a PCAP File from Scapy

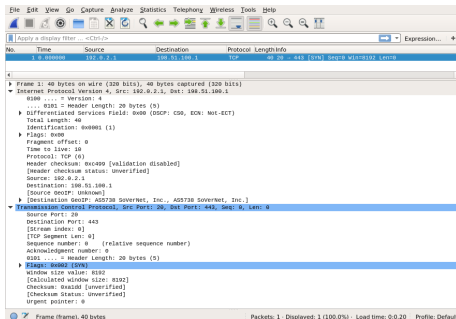


Figure: Dissection of the packet created in Scapy, in Wireshark

## Send a Packet

- The `sr1()` function sends a single packet, and returns a single packet if a reply is received.
- Start Wireshark capturing before executing the `sr1()` function.

```
1 >>> p=IP(dst="139.133.210.32")/TCP()
2 >>> a=sr1(p)
3 Begin emission:
4 . Finished sending 1 packets.
5 *
6 Received 2 packets, got 1 answers, remaining 0 packets
7 >>> a
<IP version=4 ihl=5 tos=0x0 len=44 id=0 flags=DF frag=0 ttl=47 proto=tcp checksum=0xa98d
src=139.133.210.32 dst=172.22.152.130 options=[] |<TCP sport=http dport=ftp_data
seq=3081101820 ack=1 dataofs=6 reserved=0 flags=SA window=29200 checksum=0xe9c7 urgptr
=0 options=[('MSS', 1452)] |<Padding load='v' >>>
9 >>> a.summary()
'IP / TCP 139.133.210.32: http > 172.22.152.130: ftp_data SA / Padding'
```

Listing 14: Create and Send an IP/TCP Packet

## Evil Bit

*The evil bit is a fictional **IPv4 packet header field** proposed in RFC 3514 [1], a humorous April Fools' Day RFC from 2003 authored by Steve Bellovin. The RFC recommended that the last remaining unused bit, the "Reserved Bit," in the IPv4 packet header be used to indicate whether a packet had been sent with malicious intent, thus making computer security engineering an easy problem — simply ignore any messages with the evil bit set and trust the rest.*

– Wikipedia



# Evil Bit

## Setting the Evil Bit with Scapy

- The flags in the IP header are just an attribute you can modify:

```
>>> i = IP()  
>>> i.flags = 'evil'
```

Listing 15: Setting the Evil Bit on an IPv4 Header with Scapy

# PATHspider Plugins

## ForgeSpider

```
class EvilBit(ForgeSpider, PluggableSpider):  
    2  
    name = "evilbit"  
    4    description = "Evil bit connectivity testing"  
    version = '0.0.0'  
    6    chains = [BasicChain, TCPChain, EvilChain]  
    connect_supported = ["tcpsyn"]  
    8    packets = 2  
  
    10    def forge(self, job, seq):  
        ...
```

Listing 16: Outline for Evil Bit plugin using ForgeSpider

# PATHspider Plugins

## Forging the Packets

```
1 def forge(self, job, seq):
2     sport = 0
3     while sport < 1024:
4         sport = int(RandShort())
5         l4 = (TCP(sport=sport, dport=job['dp']))
6         if ':' in job['dip']:
7             ip = IPv6(src=self.source[1], dst=job['dip'])
8         else:
9             ip = IP(src=self.source[0], dst=job['dip'])
10        if seq == 1:
11            ip.flags = 'evil'
12        return ip/l4
```

Listing 17: Creating Packets With and Without the Evil Bit

# MONROE Testbed

- MONROE's objective is to design and operate the first European transnational open platform for independent, large-scale monitoring and assessment of performance of MBB networks in heterogeneous environments
- <https://www.monroe-project.eu/>



# PATHspider on MONROE

- PATHspider is not particularly lightweight
- Return to a split model:
  - Active traffic generation and packet capture
  - Offline analysis using PATHspider

# Observations from PATHspider

- Path (e.g. source and destination IP address)
- Condition (e.g. `ecn.negotiation.succeeded`)

# Scaling Path Transparency Measurement

- with PATHspider, we've seen:
  - measurements from a single machine
  - in a single run
- How to compare measurements...
  - from multiple vantage points
  - across longer time scales?
- Answer: centralise analysis in an **observatory**
- <https://observatory.mami-project.eu/>

# Active Measurement of Path Transparency

- Methodology:
  - 1 Throw packets at the Internet
  - 2 See what happens
- Ideal: two-ended A/B testing
- Scalable: one-ended A/B testing
- Multiple sources: **isolate** on-path from near-target impairment




Figure: One-Ended vs. Two-Ended Testing



---

# Learn more about how to use PATHspider!



SIGCOMM Tutorial on  
**Repeatability and Comparability in Measurement (RCM)**  
on August 20, 2018, 2pm-5:45pm, Budapest,

- **Part I: Introduction and Topology Measurement**

- Welcome and Introduction (Brian Trammell, ETH Zurich)
- Tracebox: Topology Measurement and Impairment Discovery (Korian Edeline, U. Liege)

- **Part II: Path Transparency and Data Collection**

- PATHspider: A Tool for Controlled Hybrid Measurement (Iain R. Learmonth, U. Aberdeen)
- Observatories: Collection, Preservation, Metadata and Provenance for Active Measurement (Brian Trammell, ETH Zurich)
- The Path Transparency Observatory (Brian Trammell, ETH Zurich)

See <https://conferences.sigcomm.org/sigcomm/2018/tutorial-rcm.html>



Mirja Kühlewind: Tracing Internet Path Transparency - June 28, 2018, Vienna

16

# References I

- [1] S. Bellovin.  
The Security Flag in the IPv4 Header.  
RFC 3514 (Informational), April 2003.
- [2] Mirja Kühlewind, Sebastian Neuner, and Brian Trammell.  
On the state of ECN and TCP options on the Internet.  
In *Proceedings of the Passive and Active Measurement Conference*,  
pages 135–144, Hong Kong, China, 2013.
- [3] Iain R. Learmonth, Brian Trammell, Mirja Kühlewind, and Gorry Fairhurst.  
PATHspider: A tool for active measurement of path transparency.  
In *Proceedings of the 2016 Applied Networking Research Workshop*,  
pages 62–64, July 2016.

## References II

- [4] Brian Trammell, Mirja Kühlewind, Damiano Boppart, Iain Learmonth, Gorry Fairhurst, and Richard Scheffenegger.  
Enabling internet-wide deployment of explicit congestion notification.  
In *Proceedings of the Passive and Active Measurement Conference*,  
pages 193–205, New York, USA, 2015.