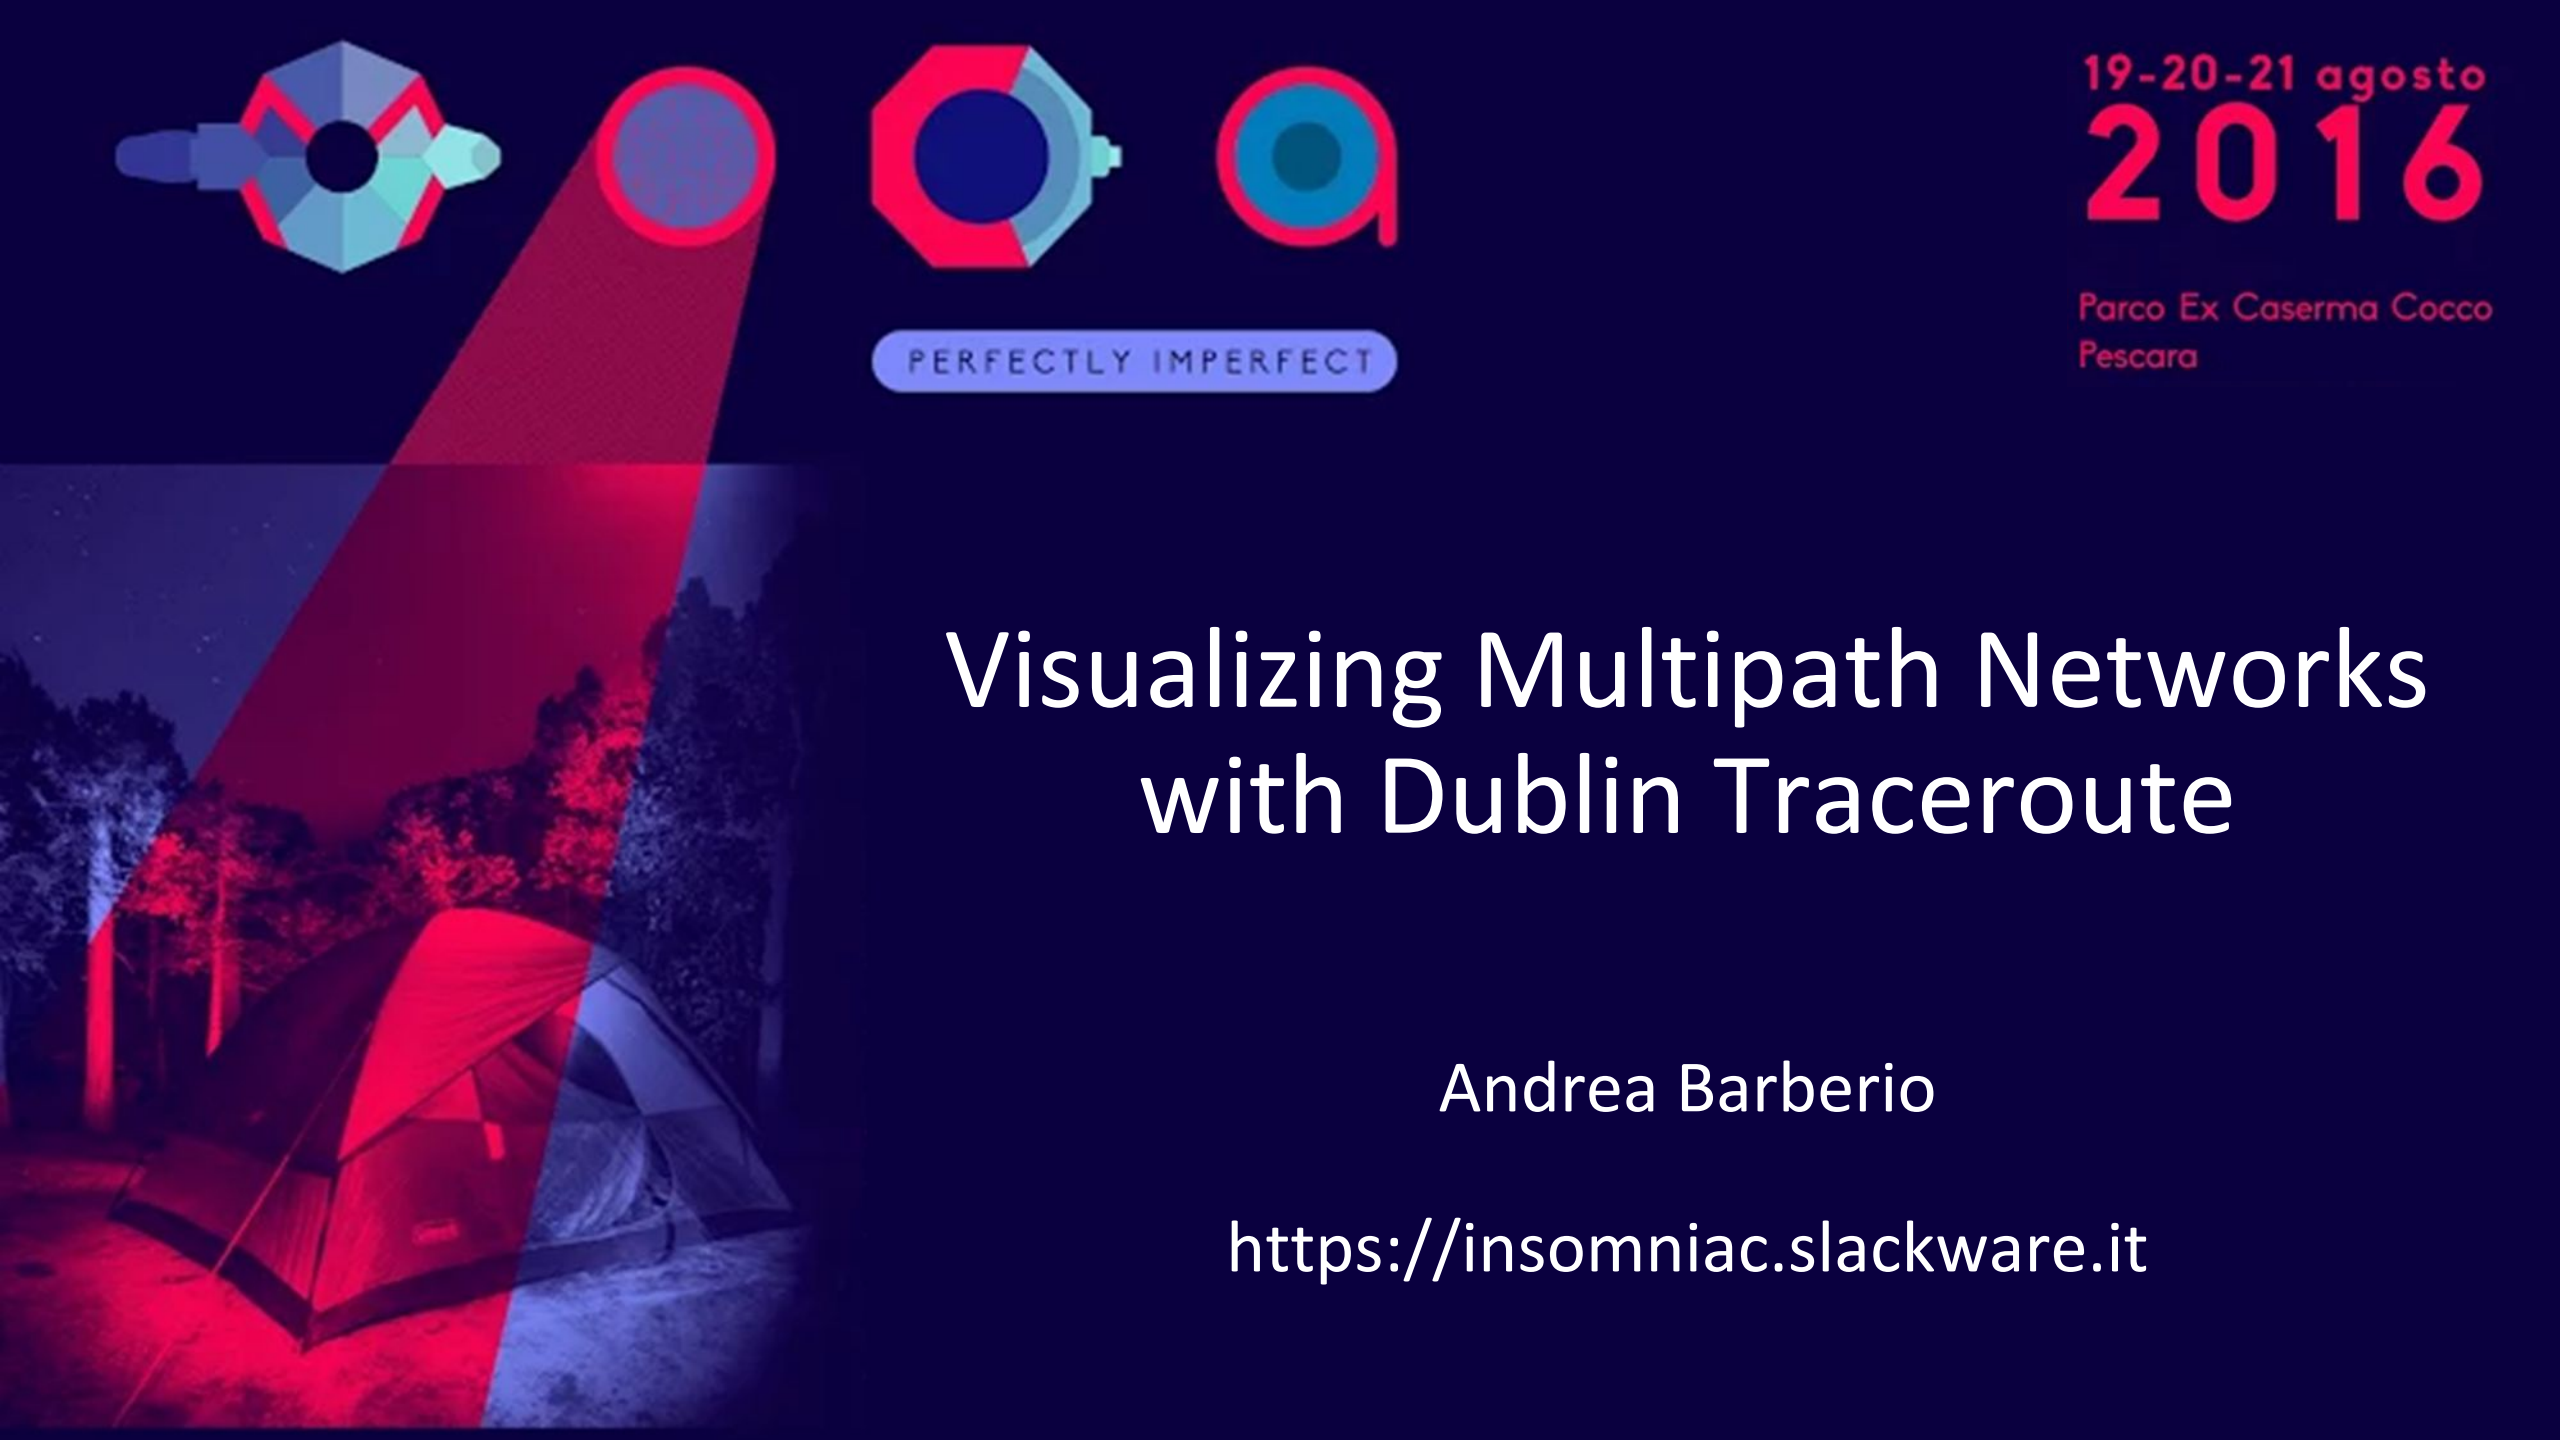

Visualizing Multipath Networks with Dublin Traceroute

Andrea Barberio

https://insomniac.slackware.it

Language - English or Italian?

Any non-italian speakers in the room?

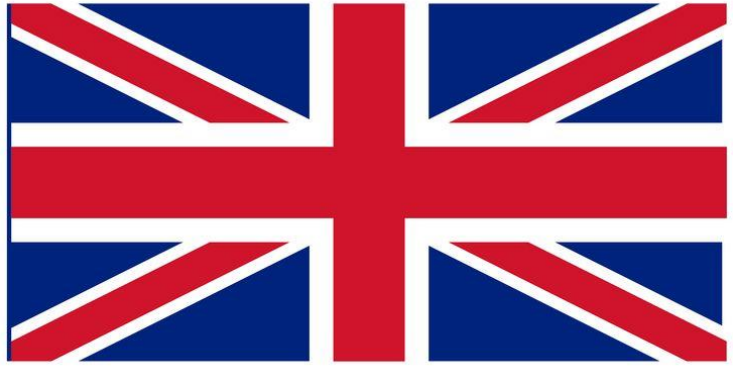

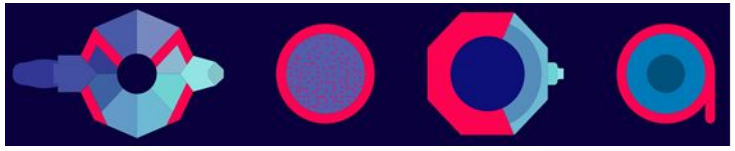

Interactivity

I would like you to be part of this talk with me

Please ask me anything at any time

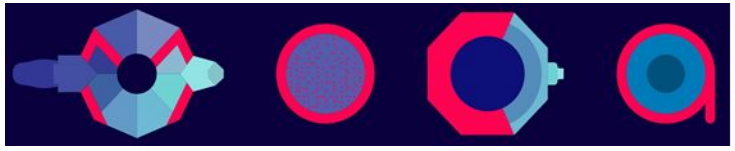

Traceroute

Probably the most widely used network tool after ping.
Send a packet with TTL=1, measure, send a packet with TTL=2, measure again, etc.

Born in 1988 (probably older than many people in this room..)

Simple concept, well understood…

...when there are no anomalies

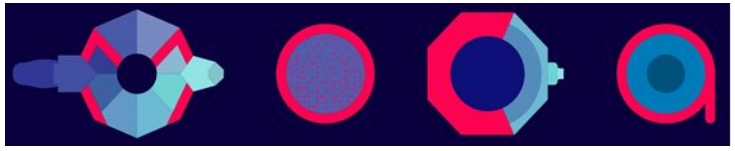

Traceroute is easy, right?

Not really. Traceroute is harder than most people believe

Anomalies are very common

Interpreting the data can be difficult

Network tomography gives only one point of view of a system, which is not necessarily correct or complete

Reconstructing the state of the network is complex and prone to failure

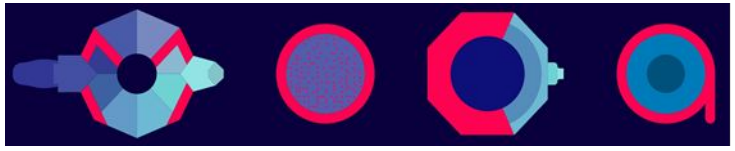

I did not convince you, right?

- MPLS tunnels
- asymmetric routes
- link aggregation ← ask me about these at the end
- anycast VIPs
- multipath load balancing
- policy changes at network boundaries
- misconfigured or misbehaving network devices
- Clos networks

What do they have in common? They are a quick way to have a bad headache

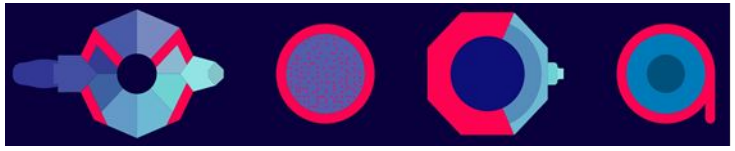

Multipath load balancing

ECMP, Equal Cost Multi-Path

Per-packet, per-flow, per-destination

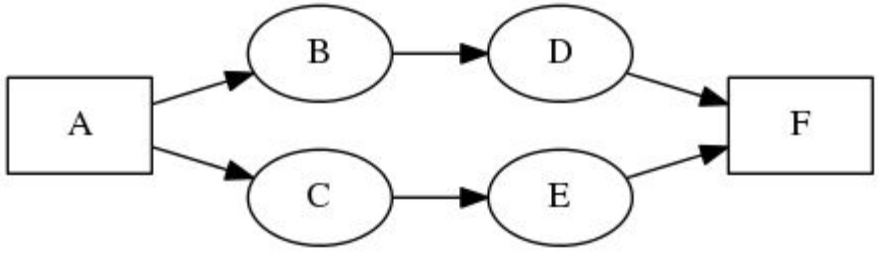

The most common is per-flow

Source: www.dublin-traceroute.net

Traceroute in ECMP environments

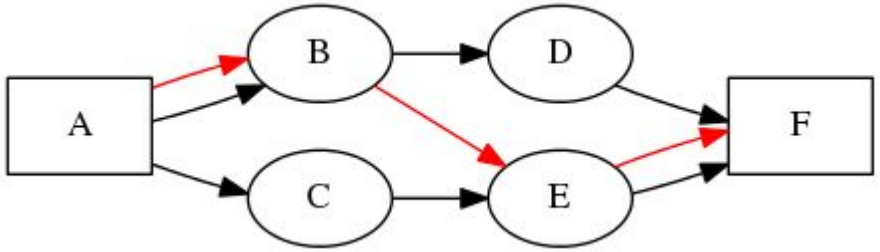

LOTS OF FUN

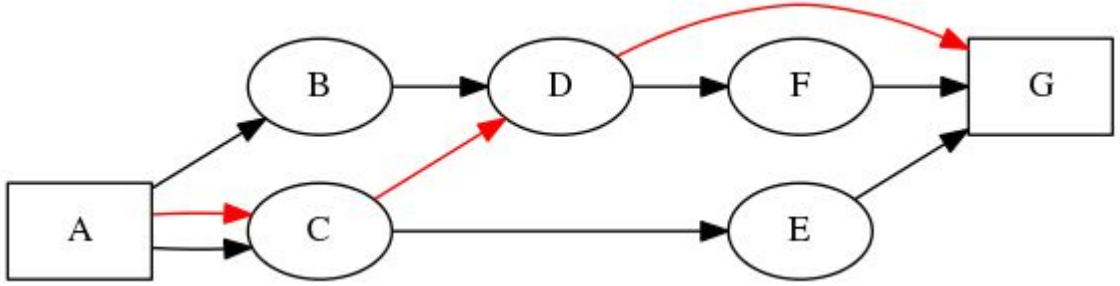

Source: www.dublin-traceroute.net

Paris-Traceroute: multipath traceroute

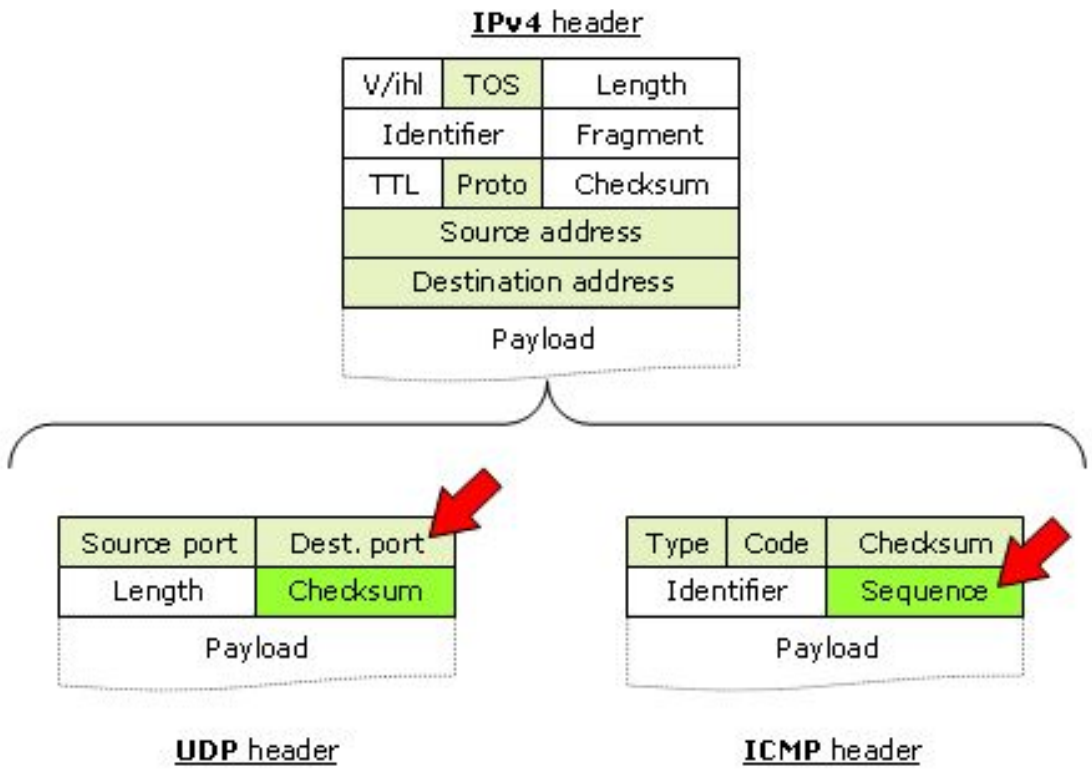

Source: www.paris-traceroute.net

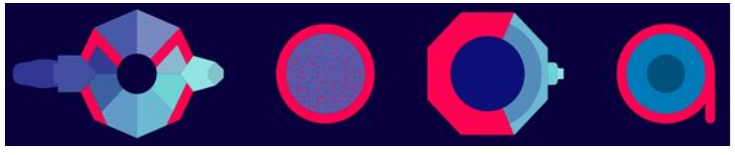

Paris-Traceroute: what's wrong with it?

Not much really, but:

- there is no NAT support (most of the current Internet)
- it's a proof-of-concept tool (not really usable in production networks)
- it's an executable only, no library (difficult to reuse)
- it's monolithic (difficult to extend)

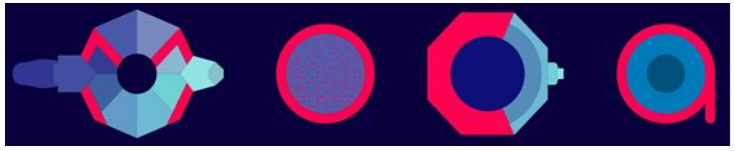

Dublin Traceroute:
NAT-aware multipath traceroute

Why?

- I wanted to do multipath traceroutes through the Internet
- I wanted a library, not just a CLI
- Easy to extend (C++11 core, Python 2/3 wrapper)
- Graphical visualization
- I heard that C++11 and Python 3 are cool
- I did it in Dublin

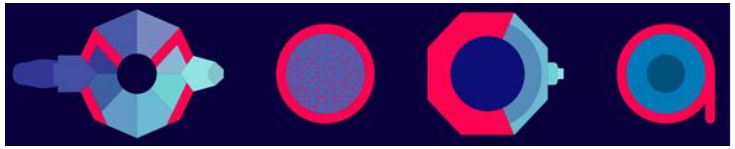

Visualizing multipath networks

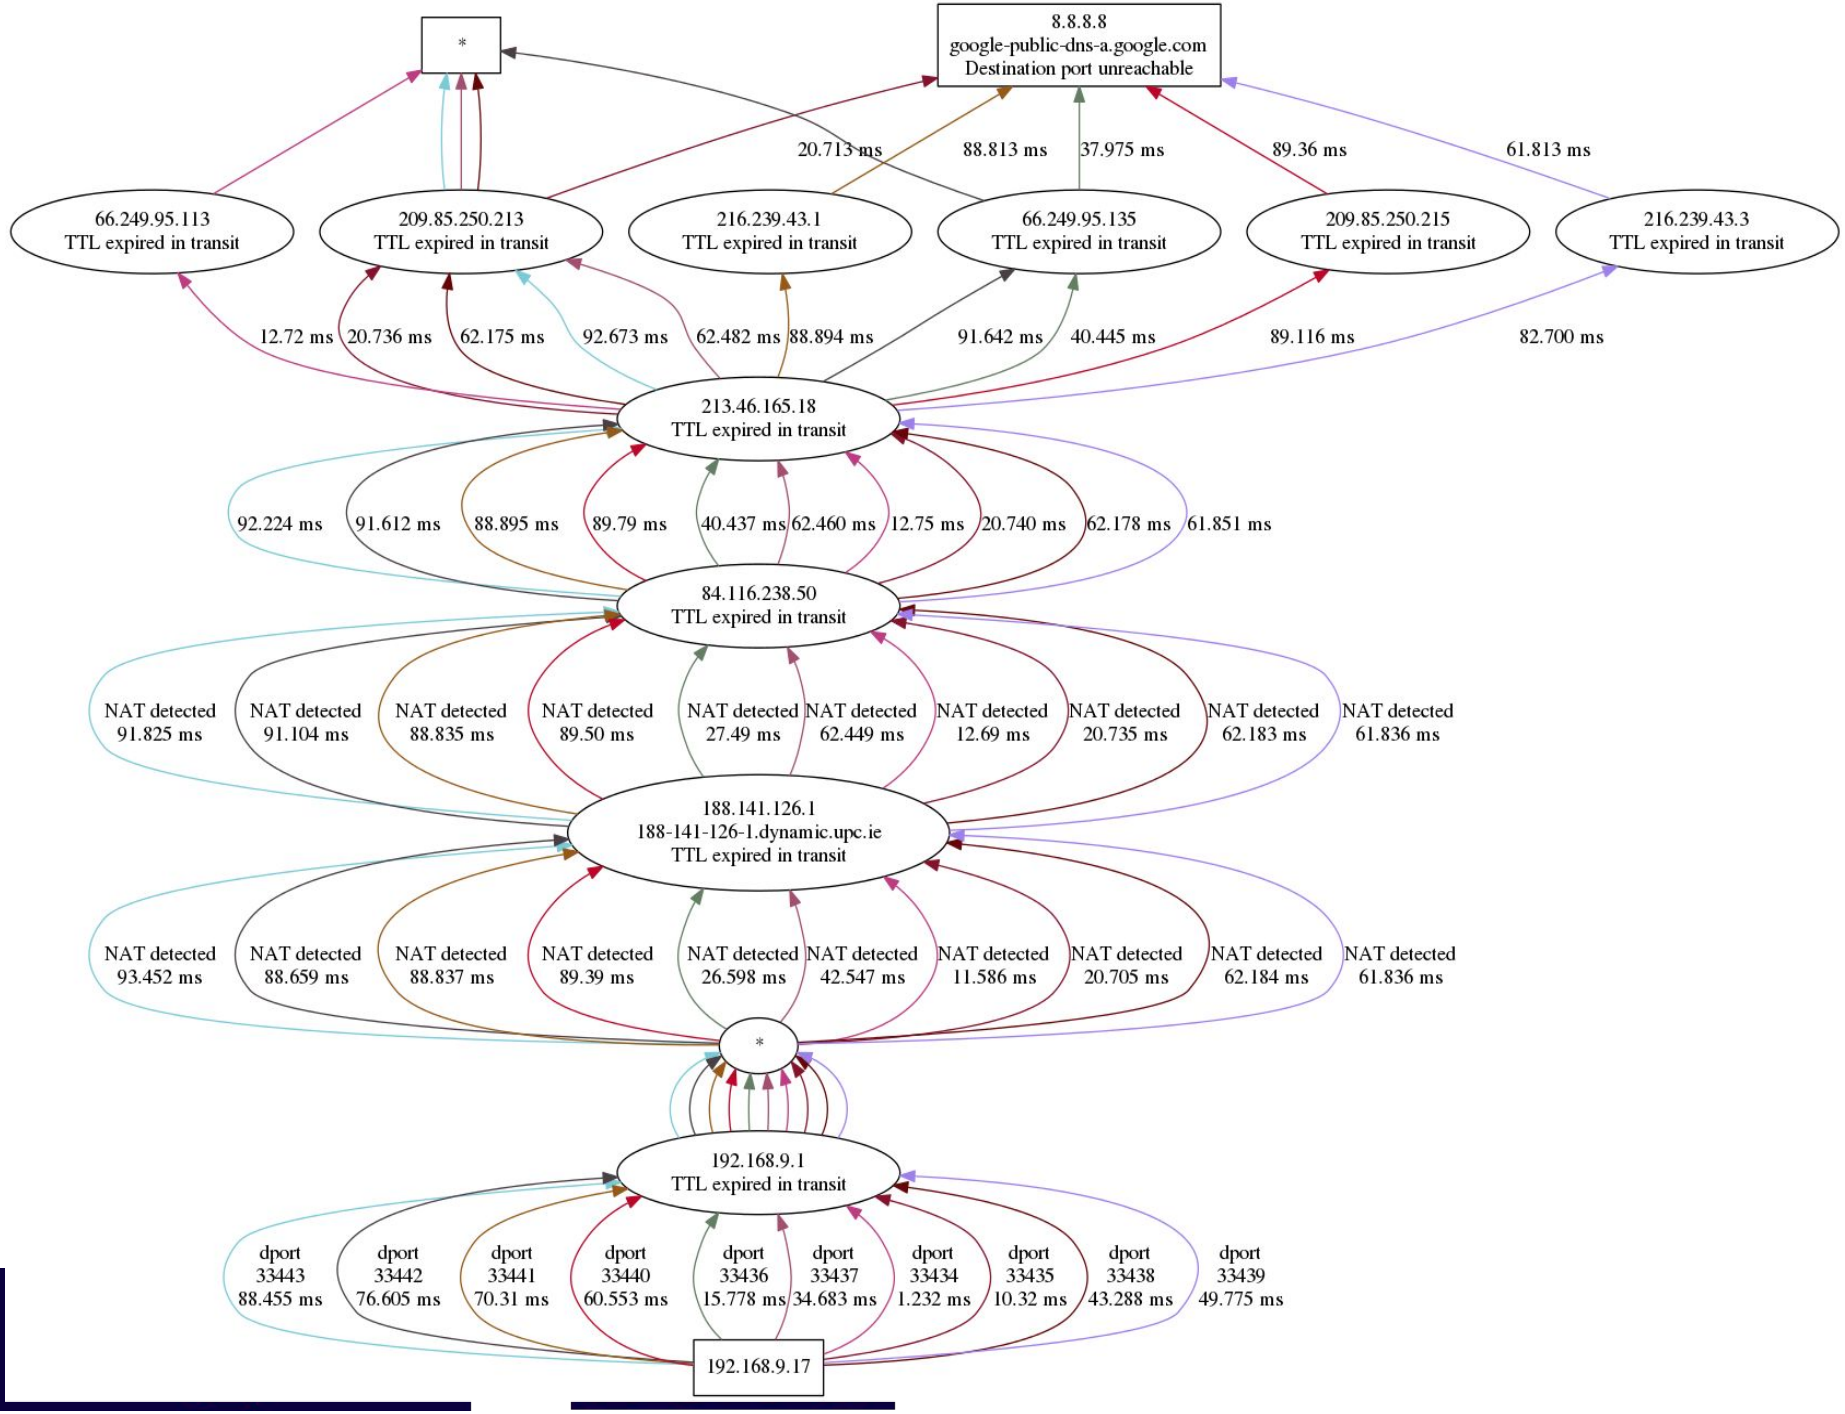

Source: www.dublin-traceroute.net

What about NAT?

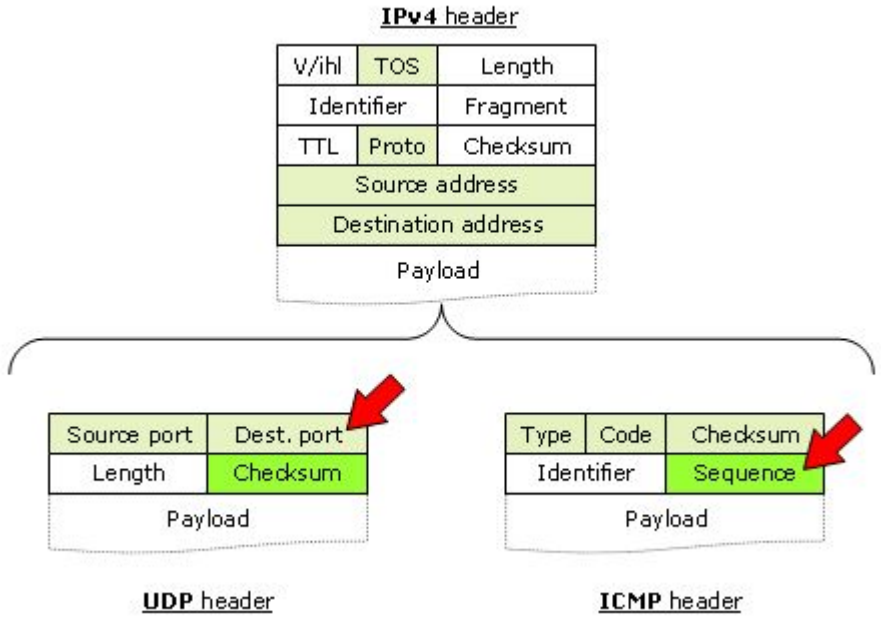

IP ID to the rescue!

The IP ID is not used for the flow hashing,
so we can manipulate it without changing path

We reuse it to identify a packet instead of the checksum for UDP. NATs manipulate the fields used for ECMP, but not the IP ID. If the checksum has changed but the ID has not, there is a NAT

We read it from the fragment returned as ICMP payload

Traceroute random fact

IANA has assigned port 33434 as base port for traceroute.

UDP and TCP have 16-bits address space for source and destination ports (65535).
2^15 (or 32768) is the first number in the second half of a 16-bits space.

2^15 + 666 = 33434

Traceroute is the tool of the beast.

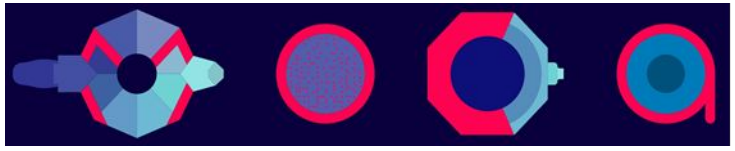

The Python library

```
from dublintraceroute import DublinTraceroute, to_graphviz

>>> dublin = DublinTraceroute(target="8.8.8.8", sport=12345, dport=33434,
...         npaths=20, max_ttl=30)
>>> results = dublin.traceroute()
>>> print(results)
>>> graph = to_graphviz(results)
>>> graph.draw("traceroute.png")
>>> graph.write("traceroute.dot")
```

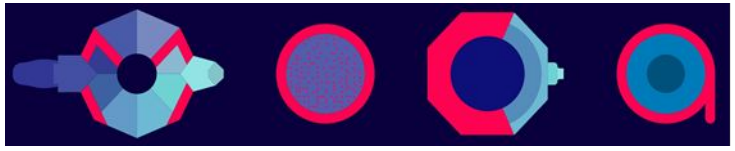

The Python interface

(install the C++ library first, via apt-get, homebrew, or sources)
pip3 install dublintraceroute

python3 -m dublintraceroute --help

python3 -m dublintraceroute 8.8.8.8
python3 -m dublintraceroute --plot trace.json

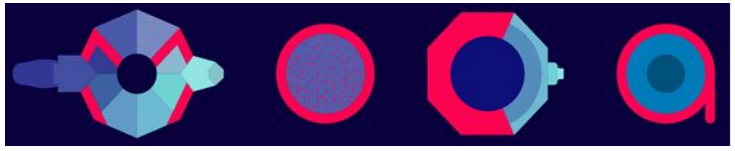

What's next?

- Support IPv6
- Identify MPLS tunnels
- Reconstruct topology and identify topology changes
- Identify censorship or two-speed Internet using multiple vantage points
- Add more probe protocols (TCP, ICMP, DNS, ..)
- Anycast VIP traversal at protocol level (e.g. DNS ChaosNet/server.id.)
- NAT64 and NAT46 detection
- Path MTU Discovery to detect latencies introduced by fragmentation
- Variable-size packets
- Unique network device identification (from distinct IPs)
- Real-time state analysis and visualization, per-link up until per-AS

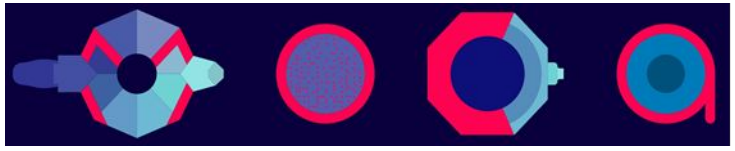

Questions?

(meanwhile.. contacts)

Andrea Barberio

insomniac@slackware.it ~ https://insomniac.slackware.it

https://dublin-traceroute.net