A quick introduction to networks

Martin L. Purschke

BROOKHAVEN
NATIONAL LABORATORY

What you should learn

- Understand an IPv4 address
- See how networks packets travel
- Understand a netmask and a broadcast address
- What a gateway is
- What is a nameserver?
- What is a time (ntp) server
- What a protocol is, and the difference between “network” and “protocol”

IPv4 and IPv6

IPv4 network was designed in the 1970's

4 billion (4*1024*1024*1024) possible addresses

That seemed a **ridiculously** large number at the time (a computer cost millions of dollars)

(Other networks, e.g. DECNet, had only 65536 (256*256) addresses)

Today, where everyone uses multiple IP addresses (laptop, iPad, smartphone), we are running out of IPv4 addresses quickly

No more IPv4 addresses in Latin America and the Caribbean

Latin America and the Caribbean have entered the IPv4 exhaustion phase; the delay in deploying Internet Protocol version 6 in our region is cause for concern.

La Casa de Internet de Latinoamérica y el Caribe, 10 June.- Today, the Internet Address Registry for Latin America and the Caribbean (LACNIC), the organization responsible for assigning Internet resources in the region, announced the exhaustion of its IPv4 address pool and expressed its concern regarding the fact that operators and governments throughout the region are delaying the deployment of Internet Protocol version 6 (IPv6).

Also, the address space is very inefficiently used

For example my institution: has 65536 addresses 15500 in use

IPv6

IPv6 network protocol was des

18446744073709551616 possible addresses (“one for each sandcorn at the beaches”)

The roll-out for the new IPv6 has been **very** slow

It is a new technology, one needs to learn

Cyber security tools (network scanners, firewalls, etc) need development

It takes a lot of infrastructure upgrades to go to IPv6 – switches, routers, manpower

Brookhaven Lab: estimated $8million

You will almost always still get a IPv4 address today.

IPv4

The Address you get is a 32bit number

Usually written as the “4-dot” notation

Eg. 196.24.32.142 (is what my Laptop has now)

Each network has a netmask and a broadcast address

inet 196.24.32.142 netmask 0xffffff00 broadcast 196.24.32.255

The broadcast and the netmask define the ”**Subnet**”

```
$ ifconfig en0
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
  ether ac:bc:32:ba:ed:6b
  inet6 fe80::1411:e20e:cf71:913d%en0 prefixlen 64 secured scopeid 0x7
  inet 196.24.32.142 netmask 0xffffff00 broadcast 196.24.32.255
```

An IPv4 subnet

A subnet is the collection of all addresses which can be reached “directly”

Subnet = all addresses where this is true:

Bit-wise “and”

Address & broadcast == Address

196 .	24	32.	142
C 4	. 1 8	. 2 0	. 8 E
1100 0100	0001 1000	0010 0000	1000 1110
	&		
1100 0100	0001 1000	0010 0000	1111 1111
	=		
1100 0100	0001 1000	0010 0000	1000 110

= 196.24.32.142 ✓

Bit-wise and:

1 0 0 1 0 0 1 1
&
0 0 0 1 1 1 1 1
=
0 0 0 1 0 0 1 1

That means that ALL addresses which start with 192.168.24 fullfill that requirement

192.168.24.78

192.168.24.199 -> 256 Addresses

Usable in practice 254 Addresses

Netmask and broadcast

Bit-wise "or" **Bit-wise "xor"**

broadcast = address | (netmask ^ 0xffffffff)

```
1111 1111 1111 1111 1111 1111 0000 0000
                    ^
1111 1111 1111 1111 1111 1111 1111 1111
                   =>
0000 0000 0000 0000 0000 0000 1111 1111
                    |
1100 0000 0001 1000 0010 0000 1000 1110
                   =>
1100 0000 0001 1000 0010 0000 1111 1111 = 192.24.32.255
```

Bit-wise xor: ("1 when they are different")

1 ^ 1 = 0
0 ^ 0 = 0
1 ^ 0 = 1
0 ^ 1 = 1

Netmask notation

netmask 0xffffff00 --- how many “leading one’s?”

`1111 1111 1111 1111 1111 1111 0000 0000`

(24) This is called a “/24” network

so you can say much shorter:

“I am on a 196.24.32.0/24 network”

What happens if I add one more “0” to my netmask? /23?

The broadcast gets “one more 1” 196.24.32.255 -> 196.24.33.255

Now everything that starts with 196.24.32 is in the subnet as before

But now also all addresses 196.24.33.X

➔ 512 addresses

Address & broadcast== address

196 . 24. 32. 142

C 4 . 1 8 . 2 0 . 8 E

1100 0100 1010 1000 0010 0000 1000 1110

&

1100 0100 1010 1000 0010 0001 1111 1111

=

1100 0100 1010 1000 0010 0000 1000 1110 = 196.24.32.142 ✓

You have now doubled the size of your subnet!

How many IP addresses does my DAQ subnet at BNL have?

```
# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:12:3F:20:7E:45
          inet addr:10.20.33.70
          Bcast:10.20.47.255
          Mask:255.255.240.0
```

If you count the bits in the netmask, you will see that this is a *huge* **/20 network**

All addresses that start with
10.20.{32,33,34,35,36,37,38,39,40,41,42,43,44,45,46,47}.XX

are part of that subnet! So it can have 4096 IP addresses.

/24	256
/23	512
/22	1024
/21	2048
/20	4096

Gateways

You can reach everything that is on the same subnet directly

So everything that starts with 196.24.32.x I can get to directly.

But that's only at most 254 addresses!

What happens if I want to access BNL's main web server,

www.bnl.gov == 130.199.3.21?

This is clearly not within my subnet range! Am I stuck?

Gateways

For anything off-subnet you go through a **gateway**

```
$ netstat -rn
Routing tables

Internet:
Destination     Gateway           Flags      Refs    Use    Netif Expire
default         196.24.32.254     UGSc         76      0      en0
. . . Lines deleted . . .
```

That gateway obviously has to be reachable on your subnet (this one is ☺)

Let's check out my BNL DAQ subnet…

```
# netstat -rn
Kernel IP routing table
Destination     Gateway         Genmask         Flags   MSS Window  irtt Iface
0.0.0.0         10.20.32.24     0.0.0.0         UG        0 0          0 eth0
. . . Lines deleted . . .
```

Remember: We calculated that all addresses that start with
10.20.{**32**,33,34,35,36,37,38,39,40,41,42,43,44,45,46,47}.**XX**
are part of my subnet. So, yes, we can reach the gateway!

So I actually did a good job! ☺

How network packets travel

If you are a car driver and want to go from one place to another, you usually have some idea how to get to your destination.

When I drive from my home to Philadelphia, I know how to get there.

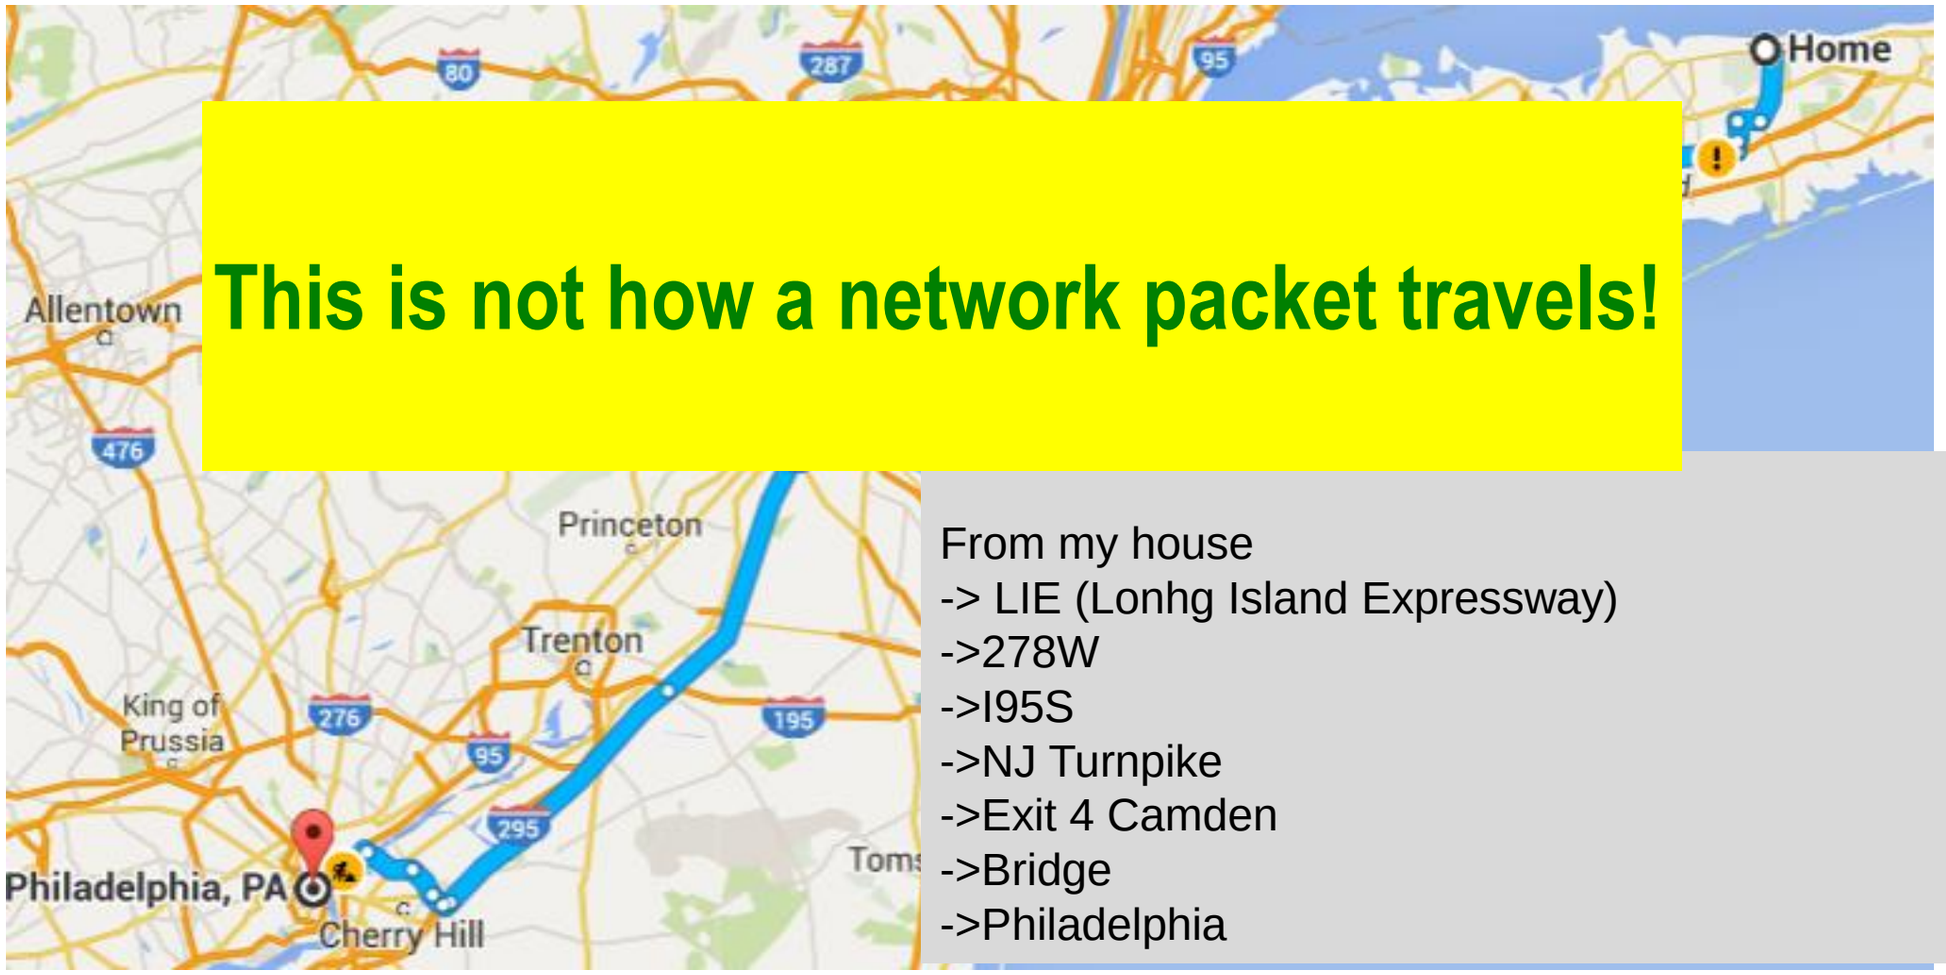

From my house
-> LIE (Lonhg Island Expressway)
->278W
->I95S
->NJ Turnpike
->Exit 4 Camden
->Bridge
->Philadelphia

How network packets travel

The network packet has NO IDEA how to get to Philadelphia

At every turn there is “someone” standing saying “Philadelphia? Make a right”. “Make a left”. And so on

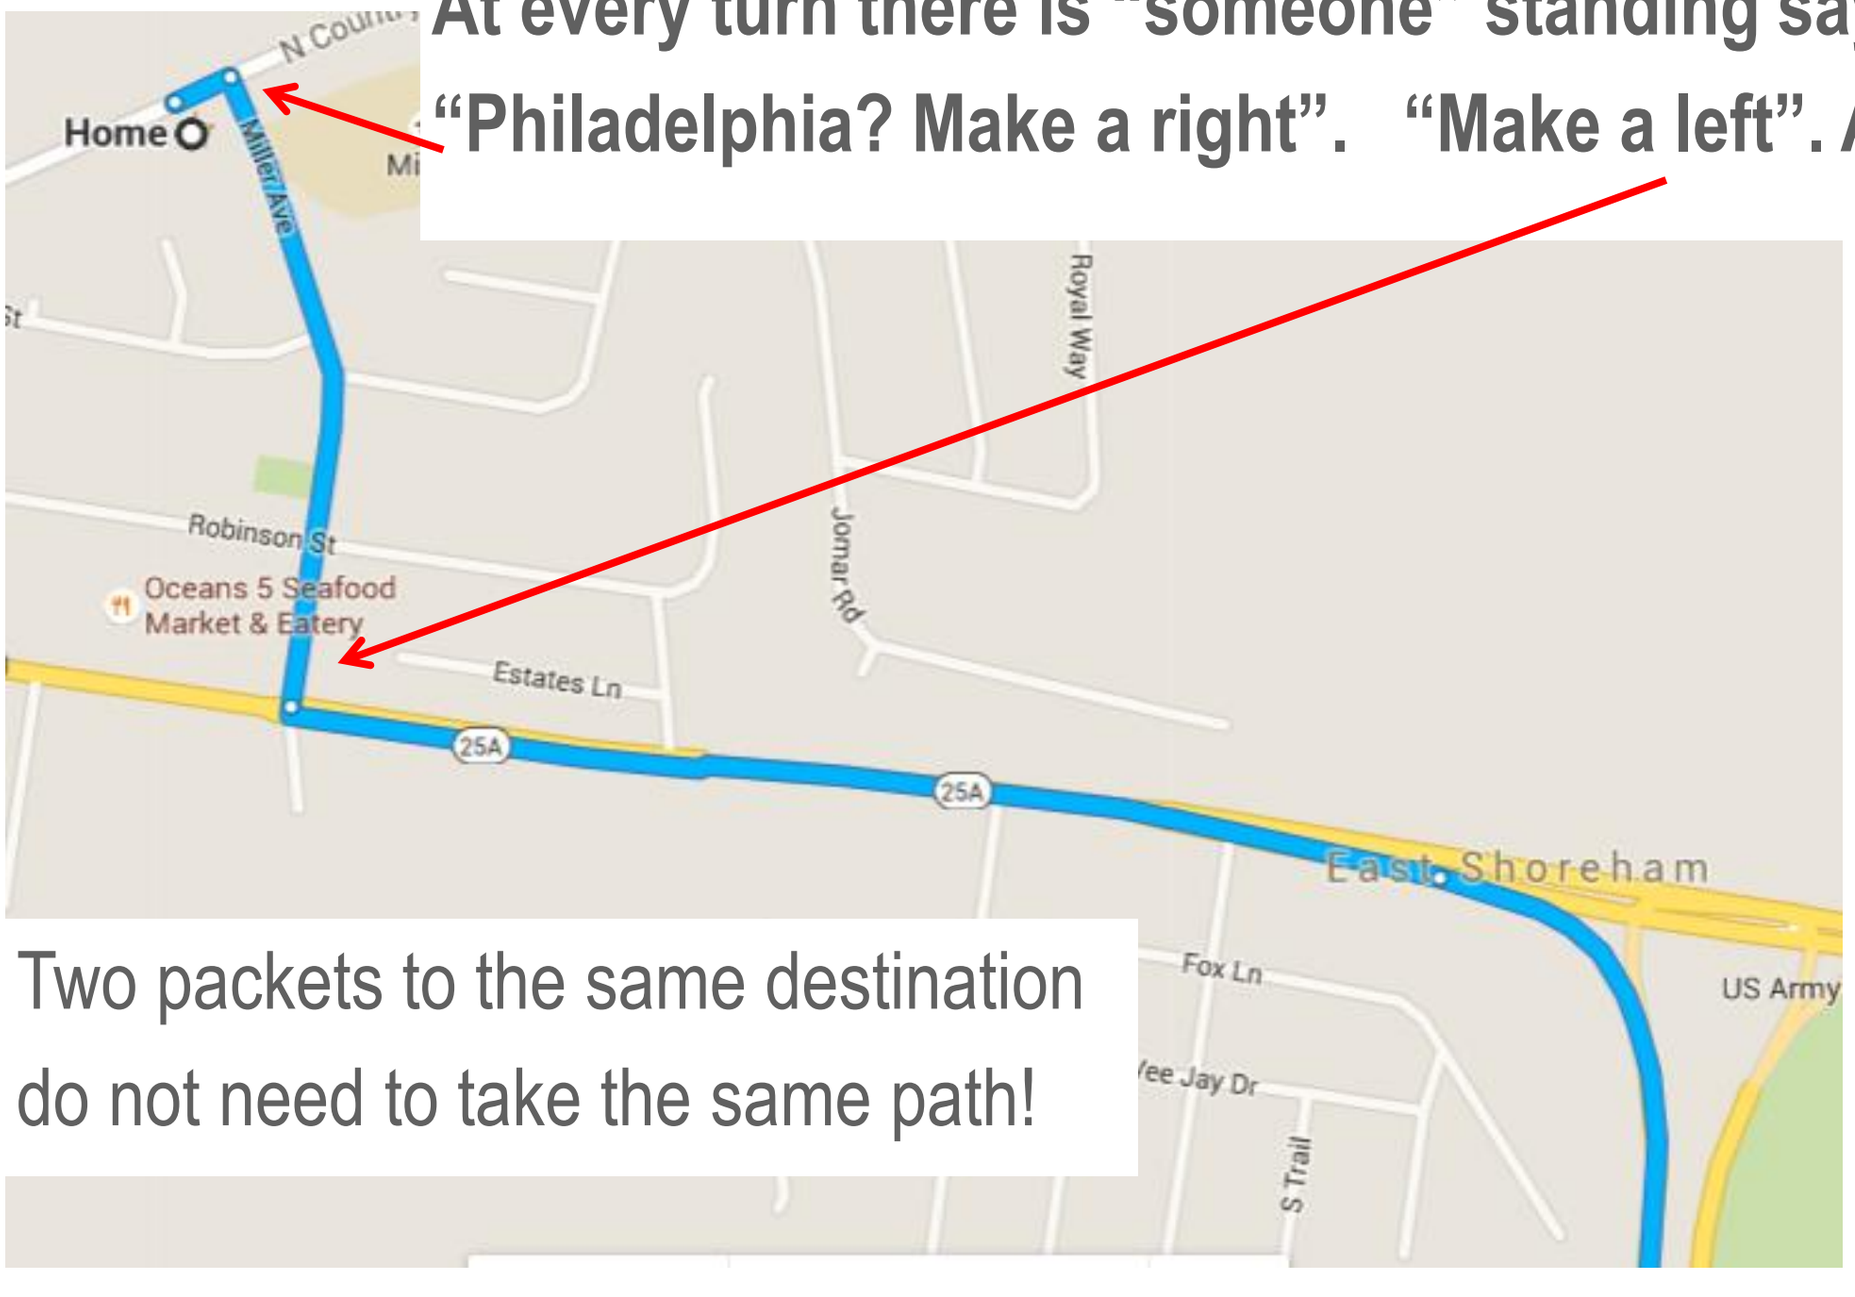

Two packets to the same destination do not need to take the same path!

How network packets travel

All it takes for a network packet is to get “off your subnet” (through the gateway)

Then your ISP/your Institute’s network people/… take over

Each time your packet arrives at a crossroads (a router), the router asks itself:

- Does this packet want to go to an IP on my subnet? (we are there!)
- Else I’ll hand it over to **my** gateway → next router, or “hop”

In this way the packet travels along a chain of routers, until it finally arrives at one that has the destination on its subnet.

We have arrived!

Going from iThemba Labs -> www.bnl.gov

```
# traceroute www.bnl.gov
traceroute to cache3.bnl.gov (130.199.3.21), 64 hops max, 52 byte packets
 1  196.24.32.254 (196.24.32.254)  8.710 ms  3.996 ms  6.004 ms
 2  sanren-ons1-vlan9.tlabs.ac.za (192.33.10.243)  10.967 ms  6.552 ms  5.597 ms
 4  ae3-0-cpt1-ir1.net.tenet.ac.za (155.232.64.74)  10.628 ms  7.205 ms  7.938 ms
 5  ae1-306-ldn1-ir1.net.tenet.ac.za (155.232.1.41)  153.431 ms  157.222 ms  145.610 ms
 6  ae1-300-ua-uk-ldn1-01.ubuntunet.net (196.32.209.221)  142.791 ms  148.375 ms  146.912 ms
 7  xe11-3-1-ua.uk.ldn1-01.ubuntunet.net (196.32.209.78)  155.786 ms  149.153 ms  149.440 ms
 8  ubuntunet.mx1.lon.uk.geant.net (62.40.124.253)  152.214 ms  143.823 ms  147.665 ms
 9  esnet-eex.lon.uk.geant.net (62.40.124.86)  194.724 ms  145.879 ms  145.822 ms
10  aofacr5-ip-c-londcr5.es.net (134.55.37.21)  305.746 ms
    newycr5-sdn-a-londcr5.es.net (134.55.220.22)  297.887 ms
    aofacr5-ip-c-londcr5.es.net (134.55.37.21)  269.949 ms
11  aofacr5-ip-a-newycr5.es.net (134.55.37.77)  219.000 ms  216.168 ms
    bnl-100ge-ip-1-aofacr5.es.net (198.124.216.10)  273.390 ms
12  * bnl-100ge-ip-1-aofacr5.es.net (198.124.216.10)  217.313 ms  327.041 ms
13  cache3.bnl.gov (130.199.3.21)  304.212 ms !Z  306.443 ms !Z *
```