

ITS 323 –IP ADDRESSING

1 IP Addresses and Dotted Decimal Notation

IP addresses are 32 bits in length. Writing and remembering 32 0's and 1's is difficult, so people have created an easier way to view IP addresses: the dotted decimal notation.

To convert a 32 bit binary IP address to its dotted decimal notation:

1. Divide the 32 bits into four 8-bit parts
2. For each part, convert the 8 bit binary number into a decimal number (between 0 and 255)
3. Combine the four decimal numbers together, separated by dots (.).

Example using the 32 bit binary address: 01110101010011001100100111110010

1. Divide into four parts: 01110101 01001100 11001001 11110010
2. Convert parts to decimal: 117 76 201 242
3. Combine the parts: 117.76.201.242

1.1 Practice Exercises

1. Convert the address to dotted decimal: 01110101001010000011101000010011
Answer: _____
2. Convert the address to binary: 53.210.12.65
Answer: _____
3. Convert the address to binary: 255.01.17.0
Answer: _____
4. Is this a valid IP address (YES or NO):
 - a. 127.255.16.0 Answer: _____
 - b. 17.0.39.256 Answer: _____
 - c. 0.0.0.0 Answer: _____

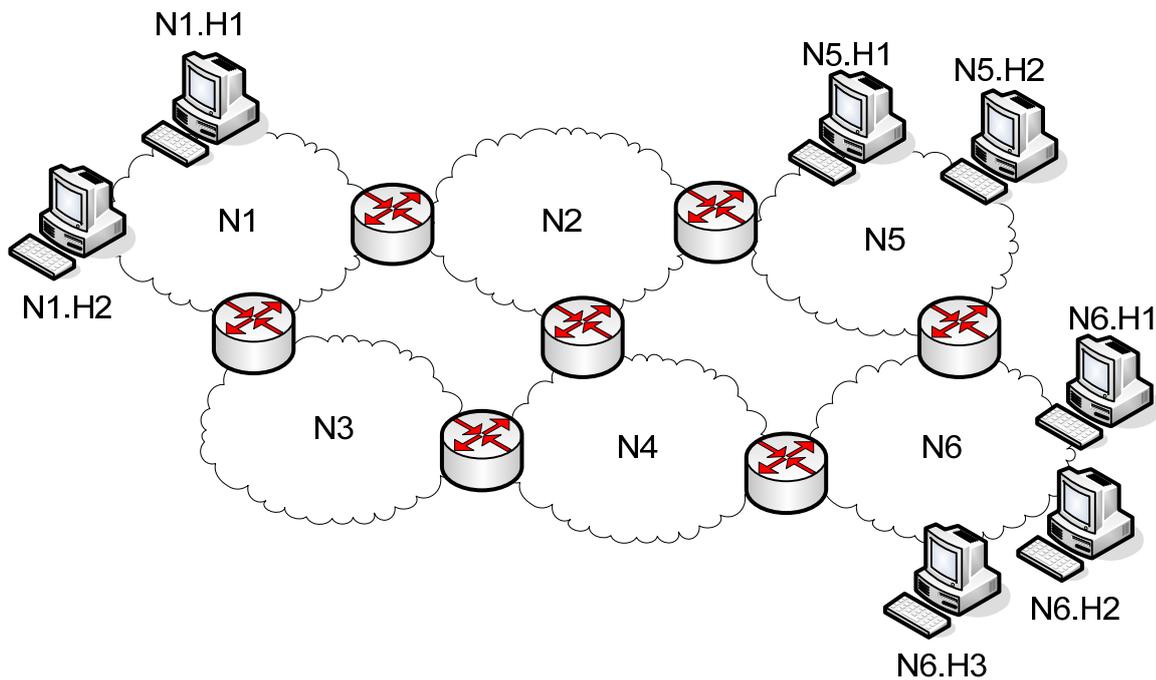
2 IP Addresses: Network + Host

IP addresses are 32 bits in length. They identify computers (at more specifically, network interfaces of computers) on the Internet. However the first part identifies a network within the Internet, and the second part identifies the computer within that network.

The figure below is a simple representation of how IP addresses identify networks and hosts. The IP addresses of computers are shown as Nx.Hy, where Nx is the part of the IP address that identifies the network the computer is attached to, and Hy is the part of the IP address that identifies the host on that network. Some conditions include:

- Every host on a network will have the same network address
- All network addresses in the internet must be unique (e.g. N1 must be different from N2)

- All host addresses in a network must be unique (e.g. on network N1, H1 must be different from H2 – BUT, H1 on network N1 does not have to be different from H1 on N5).



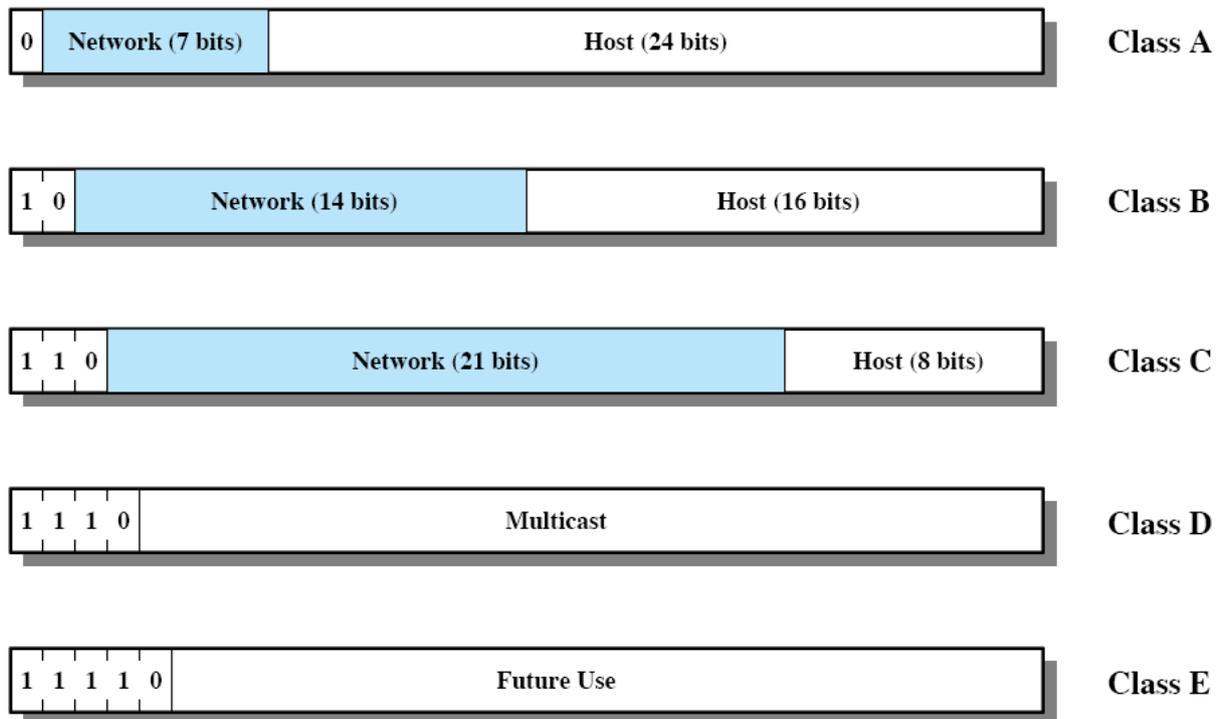
There are different schemes developed over the past 30 years to determine where to split the 32 bits into Network portion and Host portion. The general tradeoff is:

- A larger network portion allows more network on the Internet, but each network can only have a smaller number of hosts.
- A smaller network portion allows many hosts per network, but only a small number of networks.

Each scheme tries to make a tradeoff by allowing different splits between network and host portion.

3 Classful Addressing

This is the original way to split IP addresses into a Network/Host portion. The set of 2^{32} possible addresses are divided into five classes. The first three classes can be used to give an IP address to a computer – the other two classes are for special purposes, and so we will ignore them.



A computer can be given an address from any of the classes A, B or C. No matter what the class of address, any computer can communicate with any other computer (a computer with Class A address can communicate with computer with Class B address).

3.1 Class A Addresses

Class A addresses start with the first bit of 0. The next 7 bits identify the network. The remaining 24 bits identify a host. So the possible class A networks are (we will consider some exceptions later):

- 0 0000000 (dotted decimal) 0
- 0 0000001 1
- 0 0000010 2
- ...
- 0 1111110 126
- 0 1111111 127

And the possible Class A host addresses are (again, some exceptions will be consider later):

- 00000000 00000000 00000000 (dotted decimal) 0.0.0
- 00000000 00000000 00000001 0.0.1
- 00000000 00000000 00000010 0.0.2
- ...
- 11111111 11111111 11111110 255.255.254
- 11111111 11111111 11111111 255.255.255

3.2 Special Addresses

There are several IP addresses that are reserved for special purposes. That means you cannot use them to assign to a network/host. The special addresses are:

- The Class A network address that contains seven 0's: 0 0000000 (0). This is used when a computer does not have an IP address and needs one assigned.
- The Class A network address that contains seven 1's: 0 1111111 (127). This is used as a destination address if a computer wants to send to itself.
- The host address that contains all 0's. This applies to all address classes. This is used to identify a network.
- The host address that contains all 1's. This applies to all address classes. This is used as a destination if a computer wants to broadcast to all other computers on the network.

3.3 How Many Class A Networks and Hosts?

So now lets consider the Class A addresses again, and count how many class A networks are possible within an internet, and within each Class A network, how many hosts are possible.

A Class A address always starts with a 0, and the next 7 bits are the network address. Therefore there are $2^7=128$ possible network address. BUT we cannot use the network addresses that have 7 bits of 0's or 7 bits of 1's. Hence there are 126 possible Class A networks.

A Class A address has 24 bits for the host address, so there are $2^{24}=16777216$ possible host addresses. BUT we cannot use the host address of all 0's or all 1's, hence there are 16777214 possible hosts in each Class A network.

In summary, with Classful addressing, the Internet can contain a maximum of 126 Class A networks, and each of these networks can contain 16777214 hosts. Hence a maximum of $126 \times 16777214 = 2,113,928,964$ computers with Class A addresses.

But at the same time there can be Class B and Class C networks and hosts.

3.4 How Many Networks and Hosts with Classful Addresses?

As an exercise, calculate the:

1. Maximum number of Class B networks possible Answer: _____
2. Maximum number of hosts on each Class B network Answer: _____
3. Maximum number of Class C networks possible Answer: _____
4. Maximum number of hosts on each Class C network Answer: _____

Remember to consider the special host address in each class (the special network addresses only apply for Class A – you do not need to count them for class B and C).

Now calculate the total number of computers possible on the Internet if Classful addressing is used. (Hint: the answer is 4252991236 – so if you know how to calculate it correctly).

3.5 The Use of Classes

The idea behind having three classes is that there will be some organisations (no more than 126) in the Internet that need many host addresses (such as very large Internet Service Providers and Telecommunication companies). These will use Class A addresses. And then there will be many organisations that require just a few hosts (no more than 254) attached to their network (such as small business, individuals). These will use Class C addresses. And then there will be some

organisation in the middle, that need more than 254 hosts attached, but less than 65000. These will use Class B addresses.

3.6 The Problem with Classful Addressing

The problem with Classful Addressing is that most organisations in the Internet wanted Class B addresses – that is they expect to have more than 254 hosts, but certainly less than 16million hosts. And even worse, many organisations wanted many networks, e.g. Thammasat Univeristy may want a separate IP network for each faculty – that is Thammasat University takes up approximately 10 of the approximately 16,000 Class B network addresses. But many of the addresses may be “wasted”. If, for example, Thammasat needs a maximum of 1000 computers per network (10 networks in total), then it needs 10,000 IP addresses. But if Thammasat uses 10 Class B networks, it effectively uses up $10 \times 65,000 = 650,000$ IP addresses! This is a large waste.

So at that rate, there could only be about 1600 organisations in the world that could register for a Class B network. After that, there are only 126 remaining Class C networks, or many many Class A networks – but limited to 254 computers per network. It was soon realised with this scheme, the number of networks would not be enough. A new scheme was needed.

3.7 Remembering Classful Addresses

A computer can determine the class of an IP address by looking at the first bits. If the first bit is 0, Class A. If the first bit is 1 and the second bit is 0, then Class B. And if the first two bits are 1 and the third bit 0, then Class C.

But often IP addresses are in dotted decimal notation. How can you determine if an address is a Class A, B or C without converting the address to binary? Simple.

The address with the first 8 bits of 0 is a special case:

00000000 xxxxxxxx xxxxxxxx xxxxxxxx = 0.x.x.x

Class A network address is:

0 0000001 xxxxxxxx xxxxxxxx xxxxxxxx = 1.x.x.x

0 0000010 xxxxxxxx xxxxxxxx xxxxxxxx = 2.x.x.x

...

0 1111101 xxxxxxxx xxxxxxxx xxxxxxxx = 125.x.x.x

0 1111110 xxxxxxxx xxxxxxxx xxxxxxxx = 126.x.x.x

The address with first bit 0 and next 7 bits of 1 is a special case:

0 1111111 xxxxxxxx xxxxxxxx xxxxxxxx = 127.x.x.x

Class B network address is:

10 000000 00000000 xxxxxxxx xxxxxxxx = 128.0.x.x

10 000000 00000001 xxxxxxxx xxxxxxxx = 128.1.x.x

...

10 111111 11111111 xxxxxxxx xxxxxxxx = 191.255.x.x

Class C network address is:

110 00000 00000000 00000000 xxxxxxxx = 192.0.0.x

110 00000 00000000 00000001 xxxxxxxx = 192.0.1.x

...

110 11111 11111111 11111111 xxxxxxxx = 223.255.255.x

The remaining address are the special class D and E.

So in summary, the first block in dotted decimal notation tells you the class:

- Class A: 1 to 126
- Class B: 128 to 191
- Class C: 192 to 223

4 Subnet Addressing

The first improvement on Classful addressing was called Subnet addressing, or subnetting. We do not cover this in detail in the lectures (due to lack of time), but the basic concept is that the Classes are still used to identify networks, but also some bits from the Host address can be used to identify “subnets”. These subnet bits are identified using a subnet mask.

The purpose is to solve the problem like we saw with the Thammasat example. Assume Thammasat wants 10 different networks, all with a maximum of 1000 hosts. With Classful address, we saw that this used up a total of 650,000 IP addresses (when only 10,000 were really needed).

With subnet addressing, Thammasat could use one Class B address, and break that into 10 subnets. Then each subnet could have about 1000 hosts.

We will not try to describe any further, other than noting the concept of a subnet mask.

Even with subnet addressing, there remained limitations of using a fixed number of classes. Hence, the next improvement is classless addressing.

5 Classless Addressing

Classless addressing does not use the five pre-defined classes (A through to E), but instead uses an address mask to identify the split between network/host portions. This works as follows:

Given an IP address, the address mask is 32 bits that when ANDed with the IP address indicates the network portion. In other words, a 1 in the address mask indicates the corresponding bit in the IP address is part of the network portion.

The most commonly used address masks start with a sequence of 1’s and finish with a sequence of 0’s (although, in general, there can be any combination of 1’s and 0’s). An example:

IP Address: 01110101010011001100100111110010

Address mask: 11111111111111111100000000000000

Network: 011101010100110011 0000000000000000

Host: 00100111110010

In dotted decimal notation, these are:

IP Address: 117.76.201.242

Address mask: 255.255.192.0

Network: 117.76.192.0

So now, routers and computers must know about both the IP address and the address mask (in order to determine the network/host).

5.1 Short Form for Showing Address Masks

The address mask can be given in three forms:

Binary:	1111111111111111111100000000000000
Dotted Decimal:	255.255.192.0
Short-form:	/18

The short form indicates the number of 1's in the binary address mask. In our example above, there are 18 1's followed by 14 0's.

It is common to give the short form with the IP address:

117.76.201.242/18

The same special network and broadcast addresses apply.

Broadcast address:	01110101010011001111111111111111
	117.76.255.255

5.2 The Benefits of Classless Addressing

The typical approach for taking advantage of classless addressing is the assignment of addresses by ISPs. Assume an ISP has been assigned a group of addresses from the national authority. The ISP can then choose which address masks to use for different customers.

Assume there are three customers:

- Large company which has 100,000 computers
- Small company which has 800 computers
- Home office which has 5 computers

The best case, is that if the ISP could allocate 100,805 addresses in total.

For example, with a large customer the ISP may give them the address mask /15 which allows for around 130,000 hosts. But for a smaller business, the ISP may give an address mask /22 which allows for 1,000 hosts. And for home office, the ISP may give /28 which allows for 16 hosts (well, actually 14, since there are 2 special addresses).

Using classless addressing, a total of 132,112 addresses are allocated.

If classful addressing was used, the ISP only has three options: Class A, B and C. Class A is equivalent to /8, Class B is equivalent to /16 and Class C is equivalent to /24. So the large customer would have to use Class A (allows for 16million hosts), the smaller business class B (65,000) and the home office Class C (254).

Using classful addressing, a total of more than 16million addresses are allocated for these three customers, which is a significant waste of IP addresses.

Although this is an extreme example, it illustrates the benefits of classless addressing in making more efficient use of the address space.