

RCOMP - Redes de Computadores (Computer Networks)

2023/2024

Theoretical-practical lesson 02

- Networking active devices.
- Hubs and repeaters.
- Switches and bridges.
- Routers.
- LAN Ethernet networks (802.3).
- LAN wireless networks (802.11).

Networking active devices

Active network devices are all kind of devices that emit and receive signals transporting data (nodes). We can establish two categories of active devices:

- **End nodes** – they emit and receive the useful end-user's data and they are the origin, and the final destination of that data. They are workstations (clients) and the servers. A network printer or other kind of network user shared device may as well be categorized as a server.
- **Network infrastructure nodes** – they are part of the network and ensure it accomplishes its main mission: **delivering end node's data** through the network. Some of these infrastructure nodes are called **intermediate nodes**, they receive and then retransmit what they have received (they forward data). Often, infrastructure nodes use specific protocols to exchange control information between them and helping them accomplishing their missions. Under these protocols point view, these nodes act as end nodes but they don't handle end-users' data.

Yet another possible classification of active network devices relates to the OSI layer they work in. Of course, to work in a layer a device must also work on all layers below that, so we refer to the highest layer a device works with.

Network interfaces

An active network device has at least one connection to a network, it may be physical (wired network), or not physical (wireless network). Either way, the connection of a network device to a network is known as **network interface**.

Network interfaces are often called **ports**, especially for layer 1 and layer 2 intermediate nodes. Although this is not absolutely binding, often end nodes have a single network interface, and usually intermediate nodes have several network interfaces:



The lower-level components of the network interface (layer 1) emits the signal to the network medium and receive the signal from the network medium. These lower-level components of the network interface are frequently named as the transceiver (transmitter and receiver in a single word).

In a broader sense, the network interface also includes the layer 2, especially when talking about end nodes, a Network Interface Card (NIC) is an expansion card with containing a network interface:



OSI layer one intermediate nodes

Layer one specifies the physical medium, connectors, signals, line coding and modulation. The highest-level concept these devices deal with is, symbols or bits. Layer one intermediate nodes are usually named repeaters and are mainly used to amplify signals, possibly copying them from one cable to several other cables (multiple ports devices).

Amplifying a signal will increase the reach, though, when a signal is amplified, existing noise is as well amplified.

A more sophisticated repeater overcomes this issue by operating with symbols or bits, in this case when the signal is received in a port, the symbol is extracted (interpreted). Then the symbol is used to create a brand-new signal to be emitted, possibly in several output ports.


Another advantage of this type of repeater is that there's no signal type dependency between different ports. If layer 2 symbols structures (packets) are identical and data rates the same as well, these devices will be able, for instance, to interconnect a copper cable with an optical fibre cable.

When such devices have only two ports, they are usually called **repeaters**, if they have more than two ports, they are called **hubs** or **repeating hubs**.

Repeating hubs and collisions

Repeating hubs are active devices with two (in such case called repeaters) or more port cable connections. They listen for signals in every port and copy it to all ports (including the one from where it was received). The emitted signal may be an amplified version of the received signal or may be a new signal created from a received symbol.

They can even be interconnected to extend the network (image); however, they have severe limitations. The signal is spread to every network cable, so a single signal emission from any node will make the whole network busy and thus unusable for other nodes emission.



Although having a star topology, it is a shared medium network, from this point of view it is equivalent to a bus topology. At layer 2, a medium access control (MAC) protocol must be enforced because collisions will occur when two nodes emit a signal at the same time. A collision means two or more signals are mixed on the transmission medium and all data is lost.

OSI layer 2 intermediate nodes

OSI layer two introduces new concepts: the packet (PDU) and the node address. Each packet has a well-defined structure, including address information on the header (PCI). Layer two packets are usually called **frames**.

Intermediate nodes operating at layer 2, therefore retransmit frames and not signals or individual symbols. They can also use the nodes addresses information to emit frames only on the port the destination node address is connected to.

Because frames can be received and temporarily stored in memory these devices allow a full independence between what is being emitted and received in one port and what is happening in other ports. Devices ensuring these features are called **switches** (**bridges** if they have only two ports).

Upgrading from repeating hubs to switches represents a massive improvement for network performance, because each cable connection supports full duplex, all nodes can be receiving and emitting frames at the same time with no chance of collision.

Because frames can be stored in memory, the interconnection of different type cables does not even require the same data rate any longer.

The layer 2 network concept

One important concept in computer networks is the network concept itself. More precisely the layer 2 network concept, in simple words the layer 2 network embraces all nodes reachable using layer 2 frames.

From the layer 2 point of view the layer 2 network is the entire world, there is nothing beyond that, nevertheless it may be just a tiny local area network (LAN). Therefore, in layer 2 there are no network addresses, only node addresses, under the layer 2 point of view, only one network exists.

Other important concept around layer 2 networks is the **broadcast address** in local area networks. The broadcast address represents all network nodes, thus in layer 2 it represents all nodes in the layer 2 network. When a frame is sent to the layer 2 broadcast address, a copy of it will be delivered to every node in the layer 2 network, this also denotes the extent of the layer 2 network and it's called the **broadcast domain**.

At layer 3 the network concept also exists, but it is not physical related, it is logical, and thus, layer 2 networks may not precisely match layer 3 networks, although in most cases they do. Nevertheless, different layer 3 networks may share the same layer 2 network, though, **one layer 3 network cannot be dispersed throughout several layer 2 networks**.

OSI layer 3 intermediate nodes

OSI layer 3 intermediate nodes are called **routers** or gateways. Layer 3 can be viewed as a duplication of most layer 2 features: packets and node addresses. The difference is these features are now decoupled from lower layers, also the network address concept is introduced to identify different networks.

Layer 3 protocols, notably IP (IPv4 or IPv6) defines a universal packets format and uses any available layer 2 technology to transport them. However, source and destination nodes may not be connected to the same layer two network, this is where routers enter to work. Usually, a router is connected to more than one layer 2 network, it receives an IP packet transported by a frame from one layer 2 network, places the same IP packet inside a new frame of another layer 2 network and emits it.



Each layer 3 network will be identified by a network address. In IPv4 and IPv6, the network address is part of each node address, this makes routing easier because by looking at the node destination address the destination network is identified.

Unlike hubs and switches, routers can interconnect totally different layer 2 technologies, even if frame formats are different, including LAN to WAN (Wide Area Networks).

Ethernet networks – 802.3

The dominating layer 2 technology in LAN wired networks is Ethernet. It started by being a shared medium network, first using a physical bus topology with a shared coaxial cable (10base5 and 10base2) and then repeating hubs in a star topology (10baseT).

In either case MAC (Medium Access Control) is required to handle collisions.



The MAC protocol used in Ethernet is called CSMA/CD, meaning Carrier Sense Multiple Access (CSMA) with Collision Detection (CD).

CSMA means, before sending, each node must check if the transmission medium is free and if not wait until it is before sending the packet.

CD means when the node starts emitting a frame it must also listen to see if a collision is occurring, if that happens, it stops sending the frame and sends a JAM instead, informing all nodes a collision has happened, next back off and wait a period of time depending on several factors before trying again.

CSMA/CD doesn't avoid collisions, it focuses on reducing the collisions consequences by reducing each collision duration to the minimum.

CSMA/CD

As more nodes are willing to emit, collisions become more likely and thus with heavy traffic the network performance degrades significantly.

The CD proceeding also introduces some issues on frame size and network size, the problem is a collision has to be detected before the emitter finishes sending the frame, otherwise it will assume it was successful. This may not happen if the frame is too short, as it could be totally emitted (frame transmission time) before it reaches (propagation delay) the opposite network extreme where a collision will start happening.

To ensure collisions are detected in time, a compulsory minimum 64 bytes frame size is established. Depending on the transmission rate this will result in a minimum transmission time and thus a maximum network length will result to be consistent with the maximum propagation delay.

The maximum extend of an Ethernet network using CSMA/CD is called the **collision domain**, this represents the zone within the network to which collisions will be confined. If an Ethernet network uses only coaxial buses and repeating hubs the collision domain is the entire network, some Ethernet switches are capable of blocking collisions, but this is not always the case.

Of course, routers, operating at layer 3 will always block layer 2 collisions.

Ethernet node addresses

Ethernet networks have evolved in several ways, but some things were kept unchanged: **the frame format** and the **node address format**.

Each node is identified by a unique 48 bits number (6 bytes), although node addresses are required to be unique only within the layer 2 network, they are globally unique because each manufacturer is assigned a unique 24 bits number (OUI - Organizationally Unique Identifier) for the most significant bits of the node address and manages the remaining 24 bits to produce unique addresses. **Ethernet addresses**, also known as **MAC addresses** or **physical addresses** are usually represented as 6 sets of 8 bits in hexadecimal notation separated by a colon or a hyphen.

Some MAC addresses are reserved for special purposes, for instance all addresses whose first byte (left byte) has the least-significant bit with value one are multicast addresses. A frame with a multicast destination address is supposed to be retransmitted everywhere by layer 2 intermediate nodes.

The layer 2 broadcast address is made of 48 one bits, it's a special case of multicast address that represents all nodes in a layer 2 network. Thus, when a frame destination address is FF:FF:FF:FF:FF:FF all nodes within that Ethernet network will receive a copy of it.

Ethernet frame format

Even though several frame formats were initially used, the one that prevails is called **Ethernet II** or DIX (Digital Intel Xerox). It can carry up to 1500 bytes of data, this is called the **maximum transmission unit (MTU)**:

Destination node MAC address (6 bytes)	Source node MAC address (6 bytes)	Payload identifier (E-TYPE) (2 bytes)	Payload (SDU) (46 up to 1500 bytes) (MTU=1500)	Frame Check Sequence (FCS) (4 bytes)
--	---	---	---	---

Therefore, the total frame length will range from 64 up to 1518 bytes. Before sending the frame, the emitter is required to send the **Ethernet preamble** (not represented in the image above) it's made of 7 bytes with alternate zeros and ones followed by an additional byte called **start frame delimiter (SFD)**. The SFD is also made of alternate zeros and ones, but it ends with two ones. This alerts receivers of an incoming frame and allows clock synchronization. Also, the inter-packet gap (IPG) must be respected, it's the time equivalent to the transmission of 96 bits. The FCS tail is used for error detection.

By keeping the frame format and the node address format unchanged, layer 2 interoperation between different versions is guaranteed, even initial versions of Ethernet networks using coaxial cable buses can communicate with nodes connected to latest versions running at 10 Gbps over optical fiber.

Ethernet switches – no collisions

It can be said that switches saved Ethernet at a time it was becoming harder for it to compete with other technologies like Token-Ring, 100VG-AnyLAN, and especially ATM. Switches solved the Ethernet main issue: CSMA/CD.

Although externally they are like a hub, internally an Ethernet switch is a high-performance frame processing equipment, the key features are:

- At any time, it's capable of receiving a frame on any port (eventually receive different frames at the same time on all ports).
- At any time, it's capable of emitting a frame on any port (eventually emit different frames at the same time on all ports).
- It's able to temporarily store frames in memory, this is required to allow different data rates on receiving and emitting ports, it's also required when the output port is busy.
- And the fundamental feature for any switch: look at the destination node address in the frame header and retransmit the frame only on the port where such destination node is.

The three first features abolish any chance of collision, and therefore CSMA/CD is no longer required, also **because CD is not used, full-duplex is then possible.**

Ethernet switches – the MAC table

One key feature of an Ethernet switch is retransmitting frames only to the port the destination node is at. To accomplish this, the switch needs to know in which port each node is, that information is stored in the MAC table.

The MAC table is a list of known MAC node addresses, and for each the corresponding switch port. This table is dynamically managed by the switch, when the device boots it's empty, this means the switch knows nothing about surrounding node addresses, and therefore it will retransmit on all ports every frame it receives (except to the incoming port) .

However, for every frame arriving to the switch, the frame **source node** MAC address is written on the MAC table, associating to it the port where the frame was received.

Subsequently, when a frame is received the **destination node** MAC address is looked for in the MAC table, if found, the frame is retransmitted only on the associated switch port, otherwise it's retransmitted on all ports (except the incoming port). Though in case of a multicast address, the frame is always retransmitted on all existing ports (including the incoming port). Each entry in the MAC table has a time to live of only some seconds, if not refreshed before it's removed, this is important when a node is moved from one port to another.

Ethernet switches – store & forward

Totally receiving a frame in a port before starting to retransmit it on another port is called **store & forward**, this operation mode has some advantages:

- the receiving port may operate at a different data rate of the emitting port.
- collisions are not propagated because they are detected during the frame reception, therefore the receive operation will fail and there is nothing to retransmit, also if several frames are to be emitted in a single port they can wait in a queue until the port is free, so no collision will occur here as well.
- transmission errors are blocked because after receiving the entire frame error checking (FCS) by the switch is possible.

The store & forward mode has, however, one disadvantage: **latency**. There will always be a delay corresponding to the transmission time, of course this depends on the data rate and the frame size.

Other operation modes are available to overcome this, sophisticated switches can operate in several modes depending on conditions, some lower cost switches may not support store & forward and operate only in one of the alternative modes: **cut-through** or **fragment free**.

Ethernet switches - cut-through and fragment free

The **cut-through mode** provides the lowest latency. As soon as the first 6 bytes of the frame are received (destination node MAC address) the output port is immediately determined and retransmission starts, consequently, **the latency is just the time to receive 6 bytes** and not the time to receive the entire frame. Of course, data rates must be the same, collisions are propagated, also collisions may occur if the output port is busy and finally FCS errors can't be blocked.

An additional operation mode exists, it's **fragment free**. In fragment free mode the switch always receives the first 64 bytes of the frame before starting to retransmit. Due to Ethernet CSMA/CD fundamentals, all collisions must occur within these first 64 bytes, so in fragment free mode collisions are not propagated, the switch latency is now the time required to receive the first 64 bytes of the frame. However, data rates must be equal, and collisions may occur if the output port is busy. Also, FCS errors can't be blocked, that would require the entire frame to be received and often when that happens a significant part of the frame has already been transmitted.

As stated before, sophisticated switches may support several modes and apply the best one depending on the context.

Ethernet – OSI layer one

All Ethernet networks share the same layer two implementation, but they can use several different layer one types. There's a naming convention for Ethernet layer one implementations: **Data-RateSignal-TypeCable-Type**

Data-Rate – specifies the nominal data rate in Mbps, for instance **10** for 10 Mbps and **100** for 100 Mbps (Fast Ethernet), above 1000 Mbps (Gigabit Ethernet) the G letter is used for rates in Gbps, for instance **10G** for 10 Gbps.

Signal-Type – defines if a digital (**base**) or an analogic (**broad**) signal is used, current implementations only use only digital signals (**base**), the old 10broad36 implementation is now obsolete.

Cable-Type – on coaxial cable bus topologies, this is a numerical digit specifying the bus length in hundreds of meters, for instance 10base2 and 10base5 for 200 meters (180 meters in reality) and 500 meters long buses. Once the star topology was introduced, **Cable-Type** is now used to specify the cable type using letters. **T** means copper twisted pairs, **F** and **S** for multimode optical fibres and **L** for monomode optical fibres. The **X** letter is sometimes used to represent full-duplex support. In CSMA/CD mode, full-duplex is not usually supported because reception is dedicated to collision detection while emitting. Other letters and digits may be used with specific meanings.

Ethernet – physical medium

Using this naming convention, the main historical physical mediums used by Ethernet networks are:

Name	Required cable type	Description
10base5	Thick coaxial cable	Shared bus up to 500 meters long
10base2	Thin coaxial cable	Shared bus up to 180 meters long
10baseT	Copper CAT3	Uses 2 twisted pairs for 10 Mbps data rate, may run on half-duplex or full-duplex if nodes support it.
10baseF	Multimode optical fiber	Uses 2 optical fibres up to 2000 meters long in full-duplex.
100baseTX	Copper CAT5	Uses 2 twisted pairs for 100 Mbps data rate in full-duplex.
100baseFX	Multimode optical fibre	Uses 2 optical fibres up to 400 meters long in half-duplex CSMA/CD mode or up to 2000 meters in full-duplex.
100baseSX	Multimode optical fibre	Uses 2 optical fibres up to 300 meters long in full-duplex, compatible with 10baseF.

Nowadays, almost every Ethernet end node device supports rates up to 1 Gbps, called Gigabit Ethernet. When ordering new hardware, Gigabit Ethernet is currently the minimum requirement.

Ethernet – Gigabit Ethernet

Name	Required cable type	Description
1000baseT	Copper CAT5	Uses 4 twisted pairs to transmit simultaneously in both direction, supports full-duplex. Each pair is used simultaneously for sending and receiving, echo cancellation is used to remove the signal being emitted from the signal being received.
1000baseTX	Cooper CAT6	Uses 2 twisted pairs to transmit in full-duplex, not commercially implemented.
1000baseSX	Multimode optical fibre	Uses 2 optical fibres up to 550 meters long, full-duplex.
1000baseLX	Multimode or monomode optical fibre	Uses 2 optical fibres to transmit in full-duplex, up to 550 meters (multimode) or up to 2,000 meters (monomode).
1000baseLX10	Monomode optical fibre	Uses 2 optical fibres to transmit in full-duplex, up to 10,000 meters.
1000baseBX10	Monomode optical fibre	Uses a single fibre to transmit in full-duplex, up to 10,000 meters. Uses different wavelength signals for each direction.

Copper technologies (10baseT, 100baseTX, 1000baseT, and 1000baseTX) use auto-negotiation to determine the counterpart capabilities and use the best possible mode regarding the transmission rate and full-duplex support.

Ethernet – 10 Gigabit Ethernet

The most important 10 Gigabit Ethernet implementations are:

Name	Required cable type	Description
10GbaseT	Copper CAT6 or CAT6A	CAT6A or CAT7 cable is required to reach the cable maximum length of 100 meters. CAT6 cables can also be used up to 55 meters length. Uses 4 twisted pairs to transmit in both directions in full-duplex using echo cancelation.
10GbaseSR	Multimode optical fibre	Uses 2 optical fibres to transmit in full-duplex, maximum cable length depends on the fibre used, from 25 up to 400 meters.
10GbaseLR	Monomode optical fibre	Uses 2 optical fibres up to 10,000 meters long, in full-duplex.

Copper technology 10GbaseT also uses auto-negotiation, so it can be connected to 10baseT, 100baseTX, 1000baseT, and 1000baseTX counterparts.

10GbaseT was designed to be able to use older CAT6 cabling systems as far as cable length is less than 55 meters, for instance inside a datacentre.

There's a wide variety of other high-rate Ethernet physical mediums, many require special non-standard cabling systems.

Ethernet – above 10 Gbps

Higher data rate Ethernet physical mediums require special cabling systems, especially for copper cables.

Name	Required cable type	Description
40GbaseT 25GbaseT	Copper CAT8	Up to 30 meters cable length. Uses 4 twisted pairs to transmit in both directions in full-duplex.
40GbaseSR4 100GbaseSR4	Multimode optical fibre	Uses 8 optical fibres to transmit in full-duplex, maximum cable length depends on the fibre used, up to 300 meters.
40GbaseLR4 100GbaseLR4	Monomode optical fibre	Uses 8 optical fibres up to 10,000 meters long, full-duplex (four fibres for each direction).
100GbaseSR10	Multimode optical fibre	Uses 20 optical fibres, full-duplex. Maximum cable length depends on the fibre used, up to 300 meters.
100GbaseLR10	Monomode optical fibre	Uses 20 optical fibres up to 10,000 meters long, full-duplex (ten fibres for each direction).

Even higher data rates Ethernet physical mediums are to be supported, standards are being established for 200 Gbps, 400 Gbps and 1 Tbps (Terabit Ethernet). Ethernet devices working above 10 Gbps are commercially available, but still not very common.

MPO (Multi-fibre Push On) connectors

For layer one technologies requiring several optical fibres, high density connectors should be used. A MPO connector, also known as MTP (specific vendor name), can provide up to 72 optical fibre connections, but usually are used with 12 or 24 connected fibres, and thus supporting 40GbaseSR4/100GbaseSR4/40GbaseLR4/100GbaseLR4 on the first case and 100GbaseSX10/100GbaseLX10 on the second case.


Also, pre-mounted patch cords are available (images below).



Ethernet Transceiver Modules – GBIC and SFP

The name transceiver is a combination of transmitter and receiver, standing for a device capable of emitting and receiving some type of signals. Therefore, every Ethernet network interface contains an Ethernet transceiver.

However, in a stricter sense, an Ethernet transceiver module is an external device that may be connected to a special port, the advantage of such port is that by using different transceivers it can be connected to different network mediums. One of such special ports is GBIC (Gigabit Interface Converter), the images below shows a Cisco switch with two GBIC empty slots, and a couple of GBIC transceivers, one for 1000baseSX and another for 1000baseT.



Nowadays there's a more compact version of GBIC called **SFP** (small form-factor pluggable), also known as Mini-GBIC:



Ethernet Transceiver Modules – SFP+ and QSFP

Both GBIC and SFP are limited to 1 Gbps data rates, for data rates of 10 Gbps and above, other connectors are required.

SFP+ is physically similar to SFP, however, it supports data rates up to 10 Gbps. There is some compatibility between SFP+ and SFP, usually an SFP transceiver module may be plugged into an SFP+ port, of course it will operate at 1 Gbps.

Even though it's a slight deviation from the initial purpose of SFP+ ports, it is possible to directly interconnect two SFP+ ports at 10 Gbps, the following image shows such a solution with an SFP+ transceiver at each end:



For higher data rates, there's SFP28 supporting up to 25 Gbps, it's physically like SFP and SFP28 slots can take SFP+ transceiver modules.

Another option is QSFP (quad small form-factor pluggable), QSFP slots are 30% wider than SFP slots, they have four parallel data channels and can support up to 100 Gbps (QSFP28).

100GBASE-LR4 QSFP Transceiver:



Power over Ethernet (PoE)

Power over Ethernet allows passing electrical power to devices through CAT5 or better twisted-pairs copper cables. The main use for PoE is for devices that are far from telecommunications enclosures, and where there may be no power outlet available.

Some common PoE powered devices are wireless access points and video cameras. If the device supports PoE power (usually as an alternative), installing a local power outlet can then be avoided. For that purpose, a compatible power injection device must be used, this can be either an Ethernet switch with such feature, or a dedicated device usually named as **power injector** (images below).

Compatibility between the PoE powered device and the power injection device must be checked. There are several standards using different voltages, maximum current and different twisted pairs.

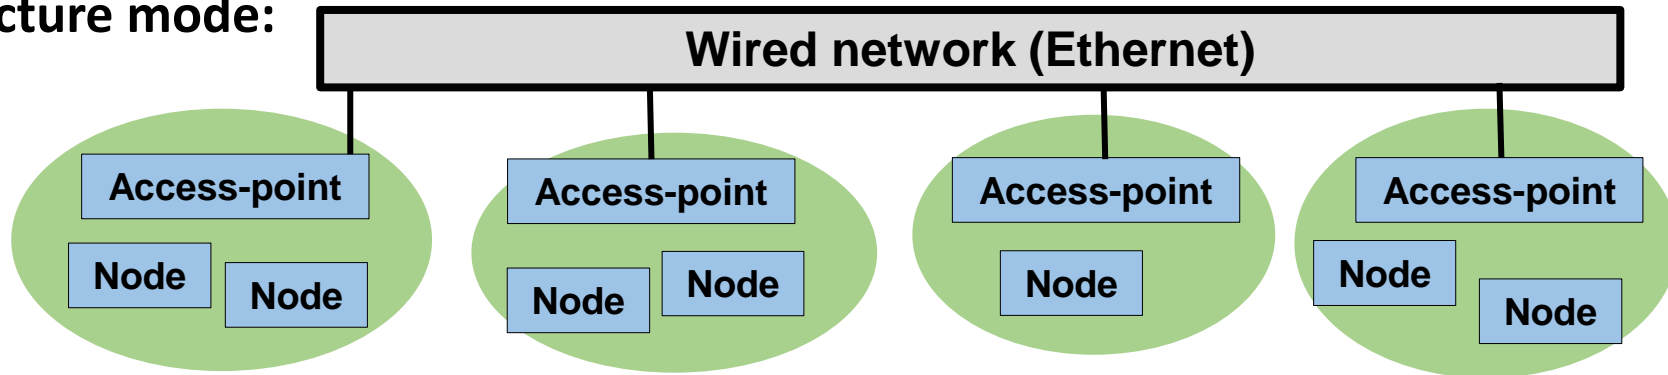


802.11 wireless local area networks (WLAN)

Also known as Wi-Fi, 802.11 is made of several standards. Alike Ethernet, it has a single layer two that can use several alternative layer one implementations, of course in this case layer one is based on radio waves.

Radio waves are analogical signals by nature, and thus, digital modulation techniques must be used to transmit data at layer one.

Although direct radio communication between end nodes is possible (ad-hoc mode), to provide **safe and solid coverage** an infrastructure of centralizing radio devices called **Wireless Access Points (WAP)** are required, this is called the **infrastructure mode**:



An access point, also known as base-station, is a layer two active device, first of all, it operates as a frame switch, retransmitting frames. Because each access point has a wired connection (Ethernet) it will not only retransmit frames between wireless nodes but also between wired nodes (Ethernet) and wireless nodes.

Wireless Access Points and Cells

In infrastructure mode, wireless nodes out of direct radio reach can nevertheless communicate with each other by using local access points which in turn communicate with each other through the wired Ethernet network. The interconnection of access points could also be wireless; however, performance would deteriorate significantly.

Each access point is responsible for managing an area of radio coverage around it called a **cell**. Due to emission power legal restrictions, the maximum effective reach of the signal is roughly less than 30 meters, thus, in practice, the cell diameter is around 50 meters.

A cell can be imagined as being a sphere centered on the access point with about 30 meters radius. This would be roughly true on empty space; however, Wi-Fi is intended for indoors space, here the scenario is rather different, microwave radio signals propagation is blocked by solid objects. Walls, beams, columns, and slabs have a significant impact on the real reach for the signal. Slabs are particularly thick and solid, and they considerably attenuate the signal passing between different floors.

When an infrastructure of access points is planned, all these propagation issues must be taken into account to attain an efficient Wi-Fi coverage of the area.

Association and roaming – BSS and ESS

When wireless workstations are within radio reach of an access-point, they may associate to the cell. Association to a cell is controlled and authorized by the access point, typically it depends on authentication, usually based on a unique pre-shared secret key (PSK) or user password authentication through central authentication services like RADIUS.

Each access point (and corresponding cell) is called a **Basic Service Set (BSS)** and is identified by a unique **Basic Service Set Identifier (BSSID)**, in fact, the BSSID is the MAC address of the access point. Cells are also identified by the **Service Set Identifier (SSID)**, this is an up to 32 bytes sequence that is usually interpreted as a human readable string. Unlike the BSSID, the SSID may not be unique for each cell. Access points that are part of the same infrastructure should all have the same SSID, thus the whole infrastructure can be seen as one by wireless workstations, this infrastructure is called an **Extended Service Set (ESS)**.


Wireless access points belonging to the same infrastructure (same SSID) allow **roaming** of clients between cells: once the client is associated to a cell, that association can be automatically and transparently transferred to a neighbour cell. Thanks to roaming, wireless clients circulating within the infrastructure are kept connected even though they are transferred from cell to cell.

802.11 - channels

It is true that different frequency radio signals don't interfere with each other, however, when data is modulated into a signal, its frequency shifts up and down, therefore adequate band wide channels are required to avoid interference between near frequency signals.

Two main frequency ranges are used in 802.11 networks: the **2.4 GHz band** and the **5 GHz band**. Each band is divided into different channels with central frequencies spaced by around 5 MHz, therefore, the band wide of these channels is 5 MHz as well.

Each layer one 802.11 standard uses these channels in different ways, for instance, **802.11g** uses channels 1 to 14 of the 2.4 GHz band, still, when 802.11g is used to modulate the signal, its frequency shifts and a **22 MHz** band wide is required, this means near channels will interfere with each other (image below).



802.11 – 5 GHz band channels

When designing an infrastructure of access points, radio coverage areas of neighbour cells will overlap, thus channels must be selected to avoid radio interference between them. On the previous image, we can see channels 1, 6 and 11 do not interfere in 802.11g, we can also see there are no non-interfering options for more than three channels.



Propagation through solid material also depends on the signal frequency, the higher the signal frequency is, the higher the attenuation will be. For instance, unlike in the 2.4 GHz band, signals on the 5 GHz band will be barely cross slabs between floors. In the 5 GHz band, cells tend to be more confined by walls and other physical obstacles.

In the 5 GHz band, the number of available channels is much higher, most of them are reserved for Dynamic Frequency Selection (DFS). DFS means channels are not manually set, instead, each access point automatically selects one to avoid overlapping with nearby detected cells.

From all this comes that selecting channels for each access point in the 5 GHz band is not an issue. On the 2.4 GHz band, however, some planning is required.

802.11 – selecting 2.4 GHz band channels

On the 2.4 GHz band, neighbour cells channels selection is particularly complex because the true is only three non-overlapping channels exist, thus avoiding interference requires a careful physical positioning of each access-point and appropriate channels selection.



Of course, with a multiple floors of a building, this turns into a three-dimensional problem.

