Introduction to WiFi Networking

Training materials for wireless trainers

The Abdus Salam
International Centre
for Theoretical Physics

United Nations
Educational, Scientific and
Cultural Organization

This 60 minute talk covers WiFi protocols, channels and radio modes, and how to use them to implement various wireless network topologies.

Version 1.7 by Rob @2010-01-10
Version 1.8 by Carlo @2010-01-11 (in Pune ;)
Version 1.9 by Ermanno @2010-01-31
Version 2.0 by Rob @2010-02-28
Version 2.1 by Carlo @2010-03-01
Version 2.2 by Rob @2010-03-01
Version 2.3 by Rob, @2010-03-12

Goals

The goal of this lecture is to introduce:

- 802.11 family of radio protocols
- 802.11 radio channels
- Wireless network topologies
- WiFi modes of operation
- Strategies for routing network traffic
- Frequently Asked Questions

The main wireless technology we will discuss is WiFi. It currently offers the greatest benefits for the lowest cost of any wireless technology. It is cheap, interoperable with equipment from different manufacturers, and can be extended to be far more than the manufacturers ever intended.

It works because WiFi uses an open standard: routers, tablet PCs, laptops, and WiFi phones can all work together because they adhere to the 802.11 standard.

Since WiFi devices comply to an open standard, the cost of equipment dropped quickly. This forced manufacturers to compete on features, rather than implement expensive and incompatible hardware with basic functionality. In the end, this was very good for consumers and manufacturers alike, as more units were shipped with a greater number of features.

ISM / UNII bands

Most commercial wireless devices (mobile phones, television, radio, etc.) use licensed radio frequencies. Large organizations pay licensing fees for the right to use those radio frequencies.

WiFi uses unlicensed spectrum. License fees are not usually required to operate WiFi equipment.

- The *Industrial, Scientific and Medical (ISM)* bands allow for unlicensed use of 2.4-2.5 GHz, 5.8 GHz, and many other (non-WiFi) frequencies.
- The *Unlicensed National Information Infrastructure (UNII)* bands allow for unlicensed use of the lower part of the 5 GHz spectrum (USA only).
- In Europe, the *European Telecommunication Standards Institute* (*ETSI*) has allocated portions of the 5 GHz band.

Note: "UNII" band is defined and regulated only in USA, other countries use different names and rules for this or similar ranges of frequencies. ISM is an international specification (ITU recommendation).
Keep in mind that ITU frequency allocations are different in different regions. Furthermore, local administrations may impose further restrictions in allowed frequencies, radio transmitted power and antenna gain.

In Europe ETSI (European Telecommunication Standards Institute) has allocated the 5470-5725 MHz band for unlicensed communication applications while in the U.S. the FCC (Federal Commission of Communications) allocated the 5725-5875 MHz band for long distance communications (maximum transmission power allowed) and the 5250-5350 MHz for medium distance. The 5150-5250 MHz is only meant for indoors communications (low power).

It is suggested to discuss the local regulation of the country in which the workshop is held.

Wireless networking protocols

The 802.11 family of radio protocols are commonly referred to as WiFi.

- ***802.11a*** supports up to 54 Mbps using the 5 GHz unlicensed bands.
- ***802.11b*** supports up to 11 Mbps using the 2.4 GHz unlicensed band.
- ***802.11g*** supports up to 54 Mbps using the 2.4 GHz unlicensed band.
- ***802.11n*** supports up to 600 Mbps using the 2.4 GHz and 5 GHz unlicensed bands.

- ***802.16*** (WiMAX) is not 802.11 WiFi! It is a completely different technology that uses a variety of licensed and unlicensed frequencies.

The specific technologies used by WiFi equipment include 802.11a, b, g, and n. 802.11n was ratified by the IEEE in September 2009, so it is a very new standard.

802.11g is backwards compatible with 802.11b, and 802.11n is backwards compatible with 802.11a when operating at 5 GHz, and b/g when operating at 2.4 GHz. By using wider channels, 802.11n loses backwards compatibility but can have much greater throughput, around 100 Mbps or more. The standard allows for even better performance by using multiple data streams. We expect equipment that supports these higher data rates to enter the market soon now that the standard has been ratified.
802.11a,b, and g are now part of the IEEE 802.11-2007 standard which encompasses all the amendments ratified up to that year, including 802.11e that deals with QoS.

Note that WiMax is not the same thing as WiFi. It is a completely different standard that operates in licensed and unlicensed frequencies.

Compatibility of standards

AP

CLIENT	802.11a	802.11b	802.11g	802.11n	802.16
802.11a	Yes			Yes @5GHz	
802.11b		Yes	Yes (slower)	Yes @2.4GHz	
802.11g		Yes (slower)	Yes	Yes @2.4GHz	
802.11n	Yes @5GHz	Yes @2.4GHz	Yes @2.4GHz	Yes	
802.16					Yes

Data rates

Note that the "data rates" quoted in the WiFi specifications refer to the raw radio symbol rate, not the actual TCP/IP throughput rate. The difference is called ***protocol overhead***, and is needed by the WiFi protocol to manage collisions, retransmissions, and general management of the link.

A good rule of thumb is to divide the radio symbol rate by two to obtain the maximum practical TCP/IP throughput. For example, a 54 Mbps 802.11a link has a maximum practical throughput of roughly 25 Mbps. An 11 Mbps 802.11b link has a maximum throughput of about 5 Mbps.

WiFI uses the same channel for both directions of traffic, so when a radio is transmitting cannot be receiving. This limits the throughput. Several vendors claim "Turbo" data rates of 108 Mbps. This is achieved by using 44MHz wide channels instead of the normal 22MHz. In most cases this is not practical.

MAC layer: CSMA vs. TDMA

802.11 WiFi uses ***Carrier Sense Multiple Access*** (***CSMA***) to avoid transmission collisions. Before a node may transmit, it must first listen for transmissions from other radios. The node may only transmit when the channel becomes idle.

Other technologies (such as WiMAX, Nstreme, and AirMAX) use ***Time Division Multiple Access*** (***TDMA***) instead. TDMA divides access to a given channel into multiple time slots, and assigns these slots to each node on the network. Each mode transmits only in its assigned slot, thereby avoiding collisions.

CSMA and TDMA are completely different media access methods. Technologies such as AirMAX or Nstreme may use 802.11 WiFi hardware, but the protocol is **not** compatible with standard WiFi!
TDMA is particularly well suited for point to point links, where there are no wasted time slots. In point to multipoint applications at short distances CSMA is more efficient.

TDMA also provides inherent Quality of Service (QoS) since the maximum time for a station to gain access to the medium is bounded.

A sort of QoS can be offered in CSMA by establishing independent queues for different type of traffic and allowing shorter interframe spacings for high priority traffic like voice, but there will be no guaranteed maximum latency.

Layer one

WiFi devices must agree on several parameters before they can communicate with each other. These parameters must be properly configured to establish "layer one" connectivity:

TCP/IP Protocol Stack	
5	Application
4	Transport
3	Internet
2	Data Link
1	Physical

- Radio channel
- Radio operating mode
- Network name
- Security features

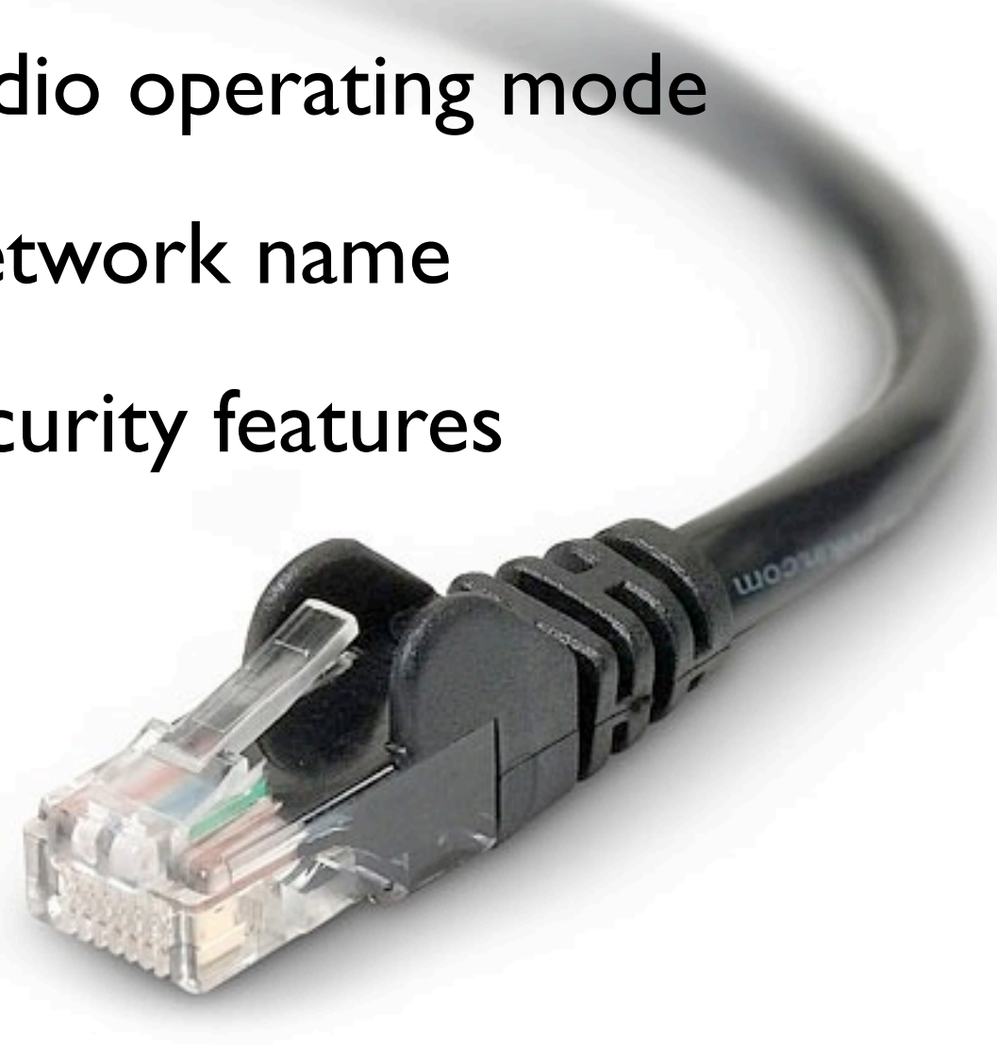

The physical layer on an Ethernet network is a wire: is it plugged in?

To establish the same layer of connectivity with WiFi, certain parameters must be set. Obviously, the devices must use the same channel or they cannot even "hear" each other. The operating mode of the radio must be properly configured or there cannot be communication. The name of the network (also called the ESSID) must match. And any security features must also be configured correctly.

Unless these parameters are properly set, it is as if the "cable" is not yet plugged in. This configuration procedure will be analyzed with greater details in another lecture (AP configuration).

802.11 WiFi Channels

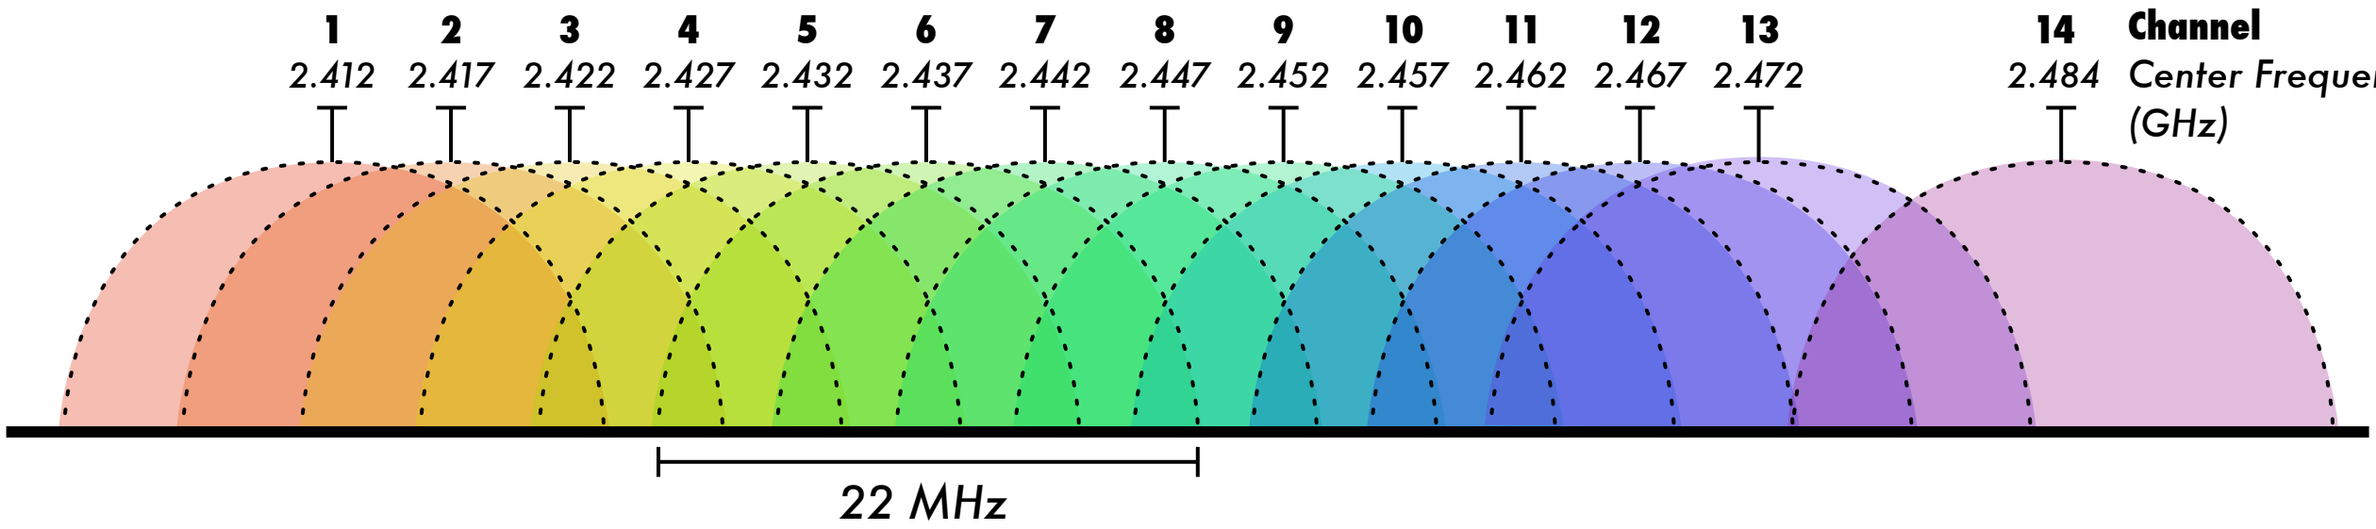

WiFi devices must use the same channel in order to communicate with each other. They send and receive on the same channel, so only one device may transmit at any time. This kind of connection is called ***half-duplex***.

In half-duplex communications, only one device may transmit at any time. This may very different from Ethernet networks, where full duplex mode (simultaneous transmit and receive) is also allowed by the standard for some hardware configurations. As we will see, this becomes an important consideration when making long distance wireless networks.

Non-overlapping channels: 1, 6, 11

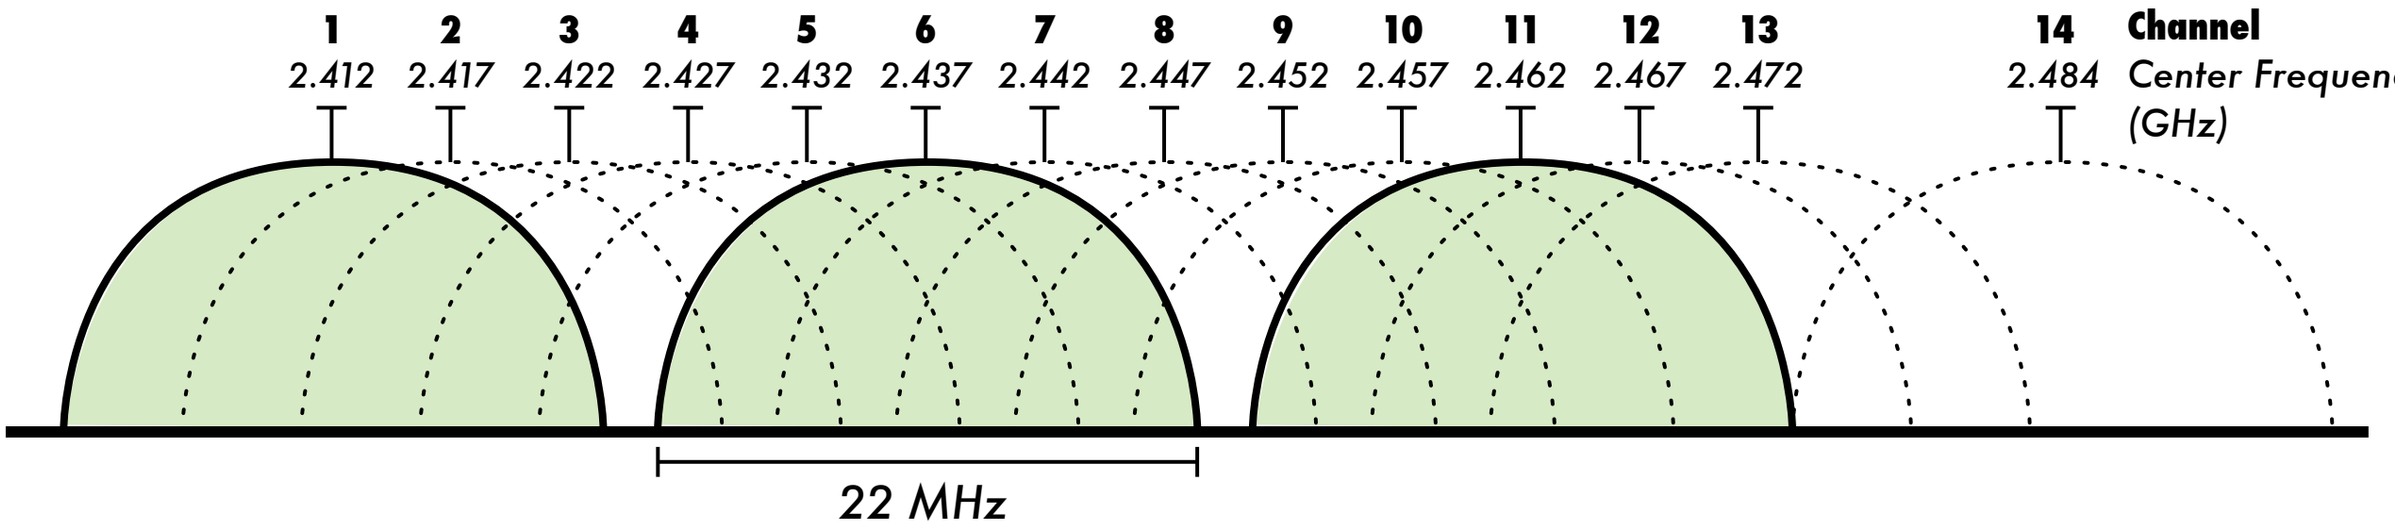

You can choose to use channels that do not overlap at all. If you operate 802.11 equipment on non-overlapping channels, they will not interfere with each other.

For example, channels 1, 6, and 11 do not overlap with each other.

(This has been explained with greater details and more examples in the lecture on "Comparative use of unlicensed spectrum")

AP channel re-use

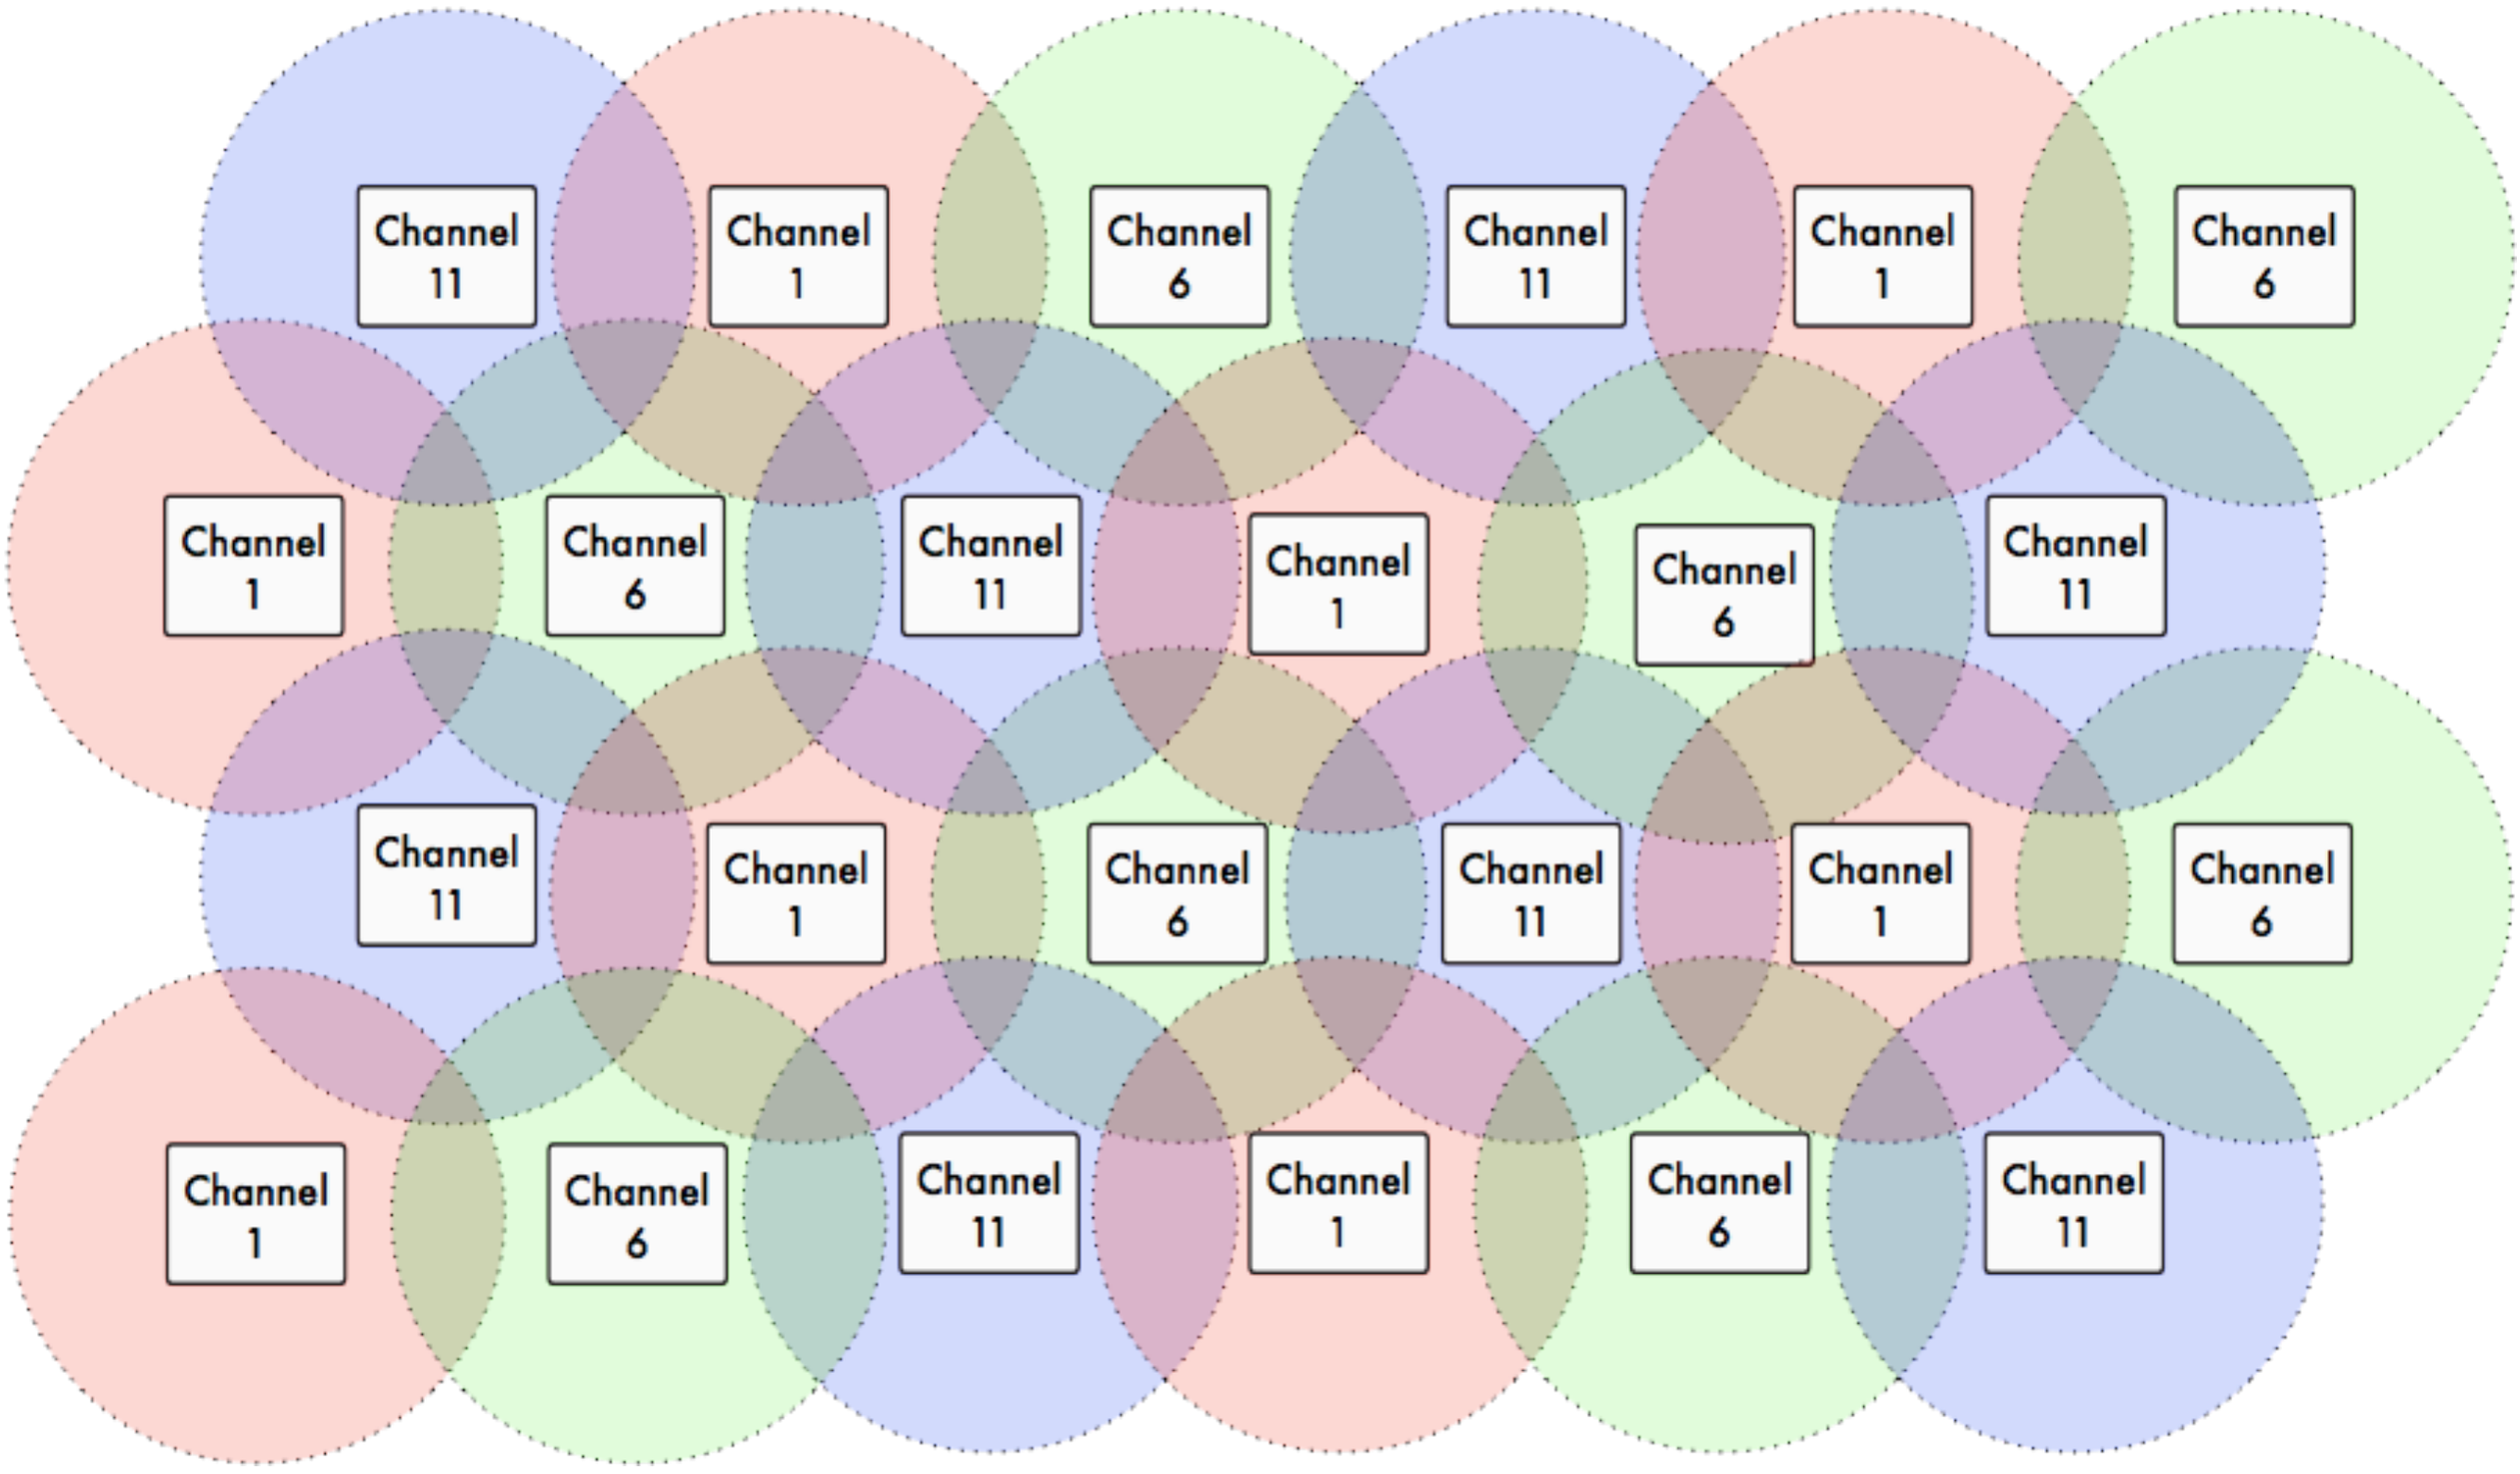

This picture represents one scheme for selecting channels for your access points that do not overlap. If you are careful about AP placement, you can cover an arbitrarily large campus using only three channels, with no adjacent channel interference.
Of course, the real world is never this pretty, so this is an ideal diagram. The coverage of access points is never a perfect circle. Also, consider the topographical problem of extending the network in three dimensions, as you have in multiple floors of a building.

Wireless network topologies

Any complex wireless network can be thought of as a combination of one or more of these types of connections:

- ***Point-to-Point***
- ***Point-to-Multipoint***
- ***Multipoint-to-Multipoint***

To step back for a moment from WiFi, let's consider a network using any generic wireless technology. These are the three fundamental topologies that can be combined to create an arbitrarily complex wireless network.

It's important to come back to these basic building blocks. As your network grows, it is easy for the complexity of the network to run away with you. But if you reduce parts of the network to just one of these three topologies, it becomes clear how information will flow through the network.

Keep in mind that none of these topologies is the "best" kind of network. Each has strengths and weaknesses, and should be applied appropriately to the networking problem you need to solve.

Point to Point

The simplest connection is the ***point-to-point*** link.

These links can be used to extend a network over great distances.

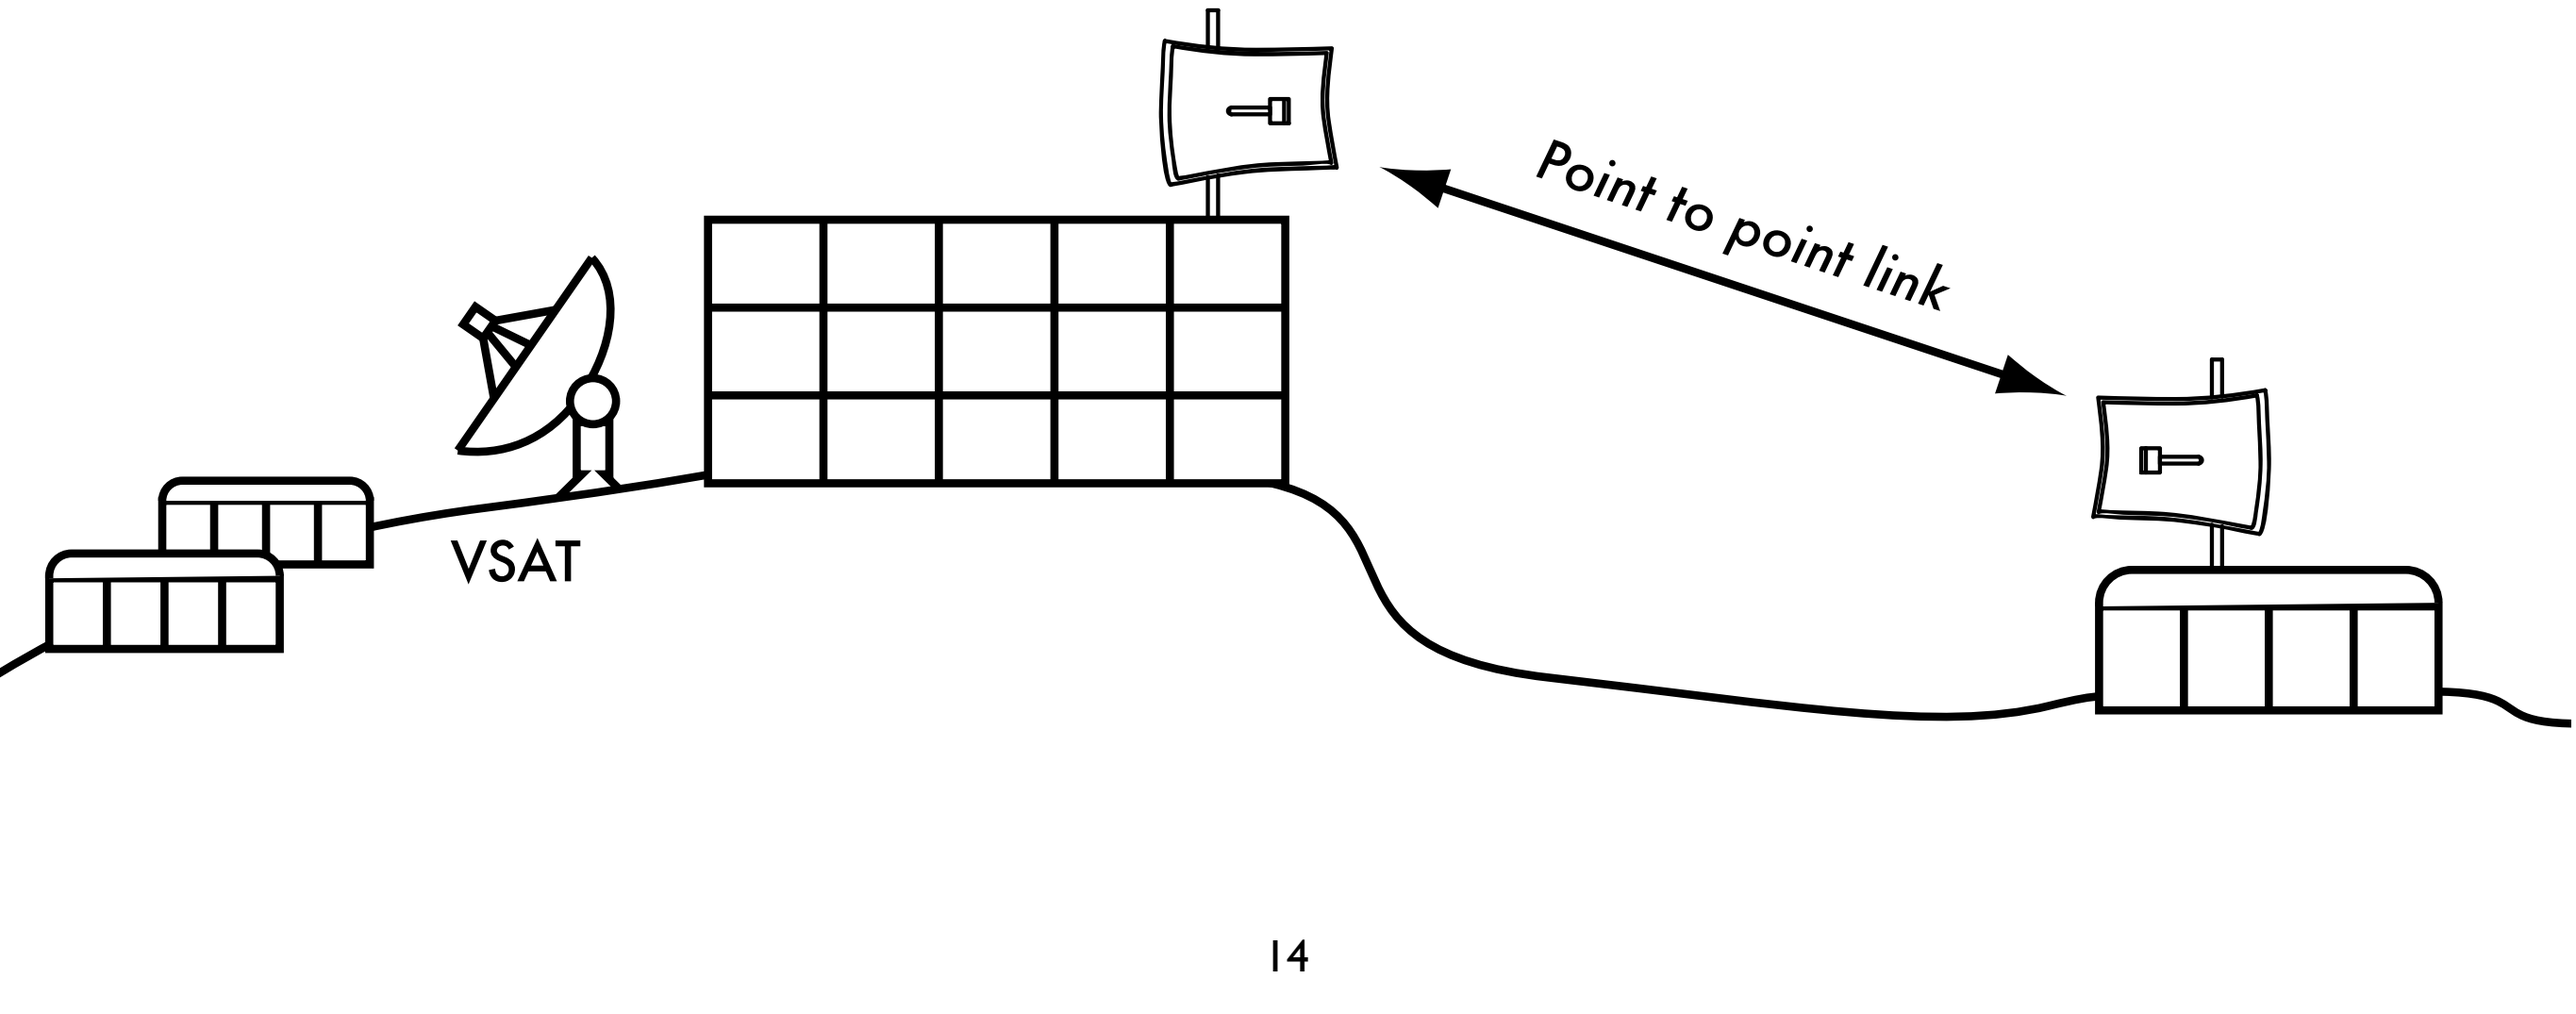

Point to point links have the greatest possible throughput out of any of the basic network topologies, since there is very little contention for resources.

Point to Multipoint

When more than one node communicates with a central point, this is a ***point-to-multipoint*** network.

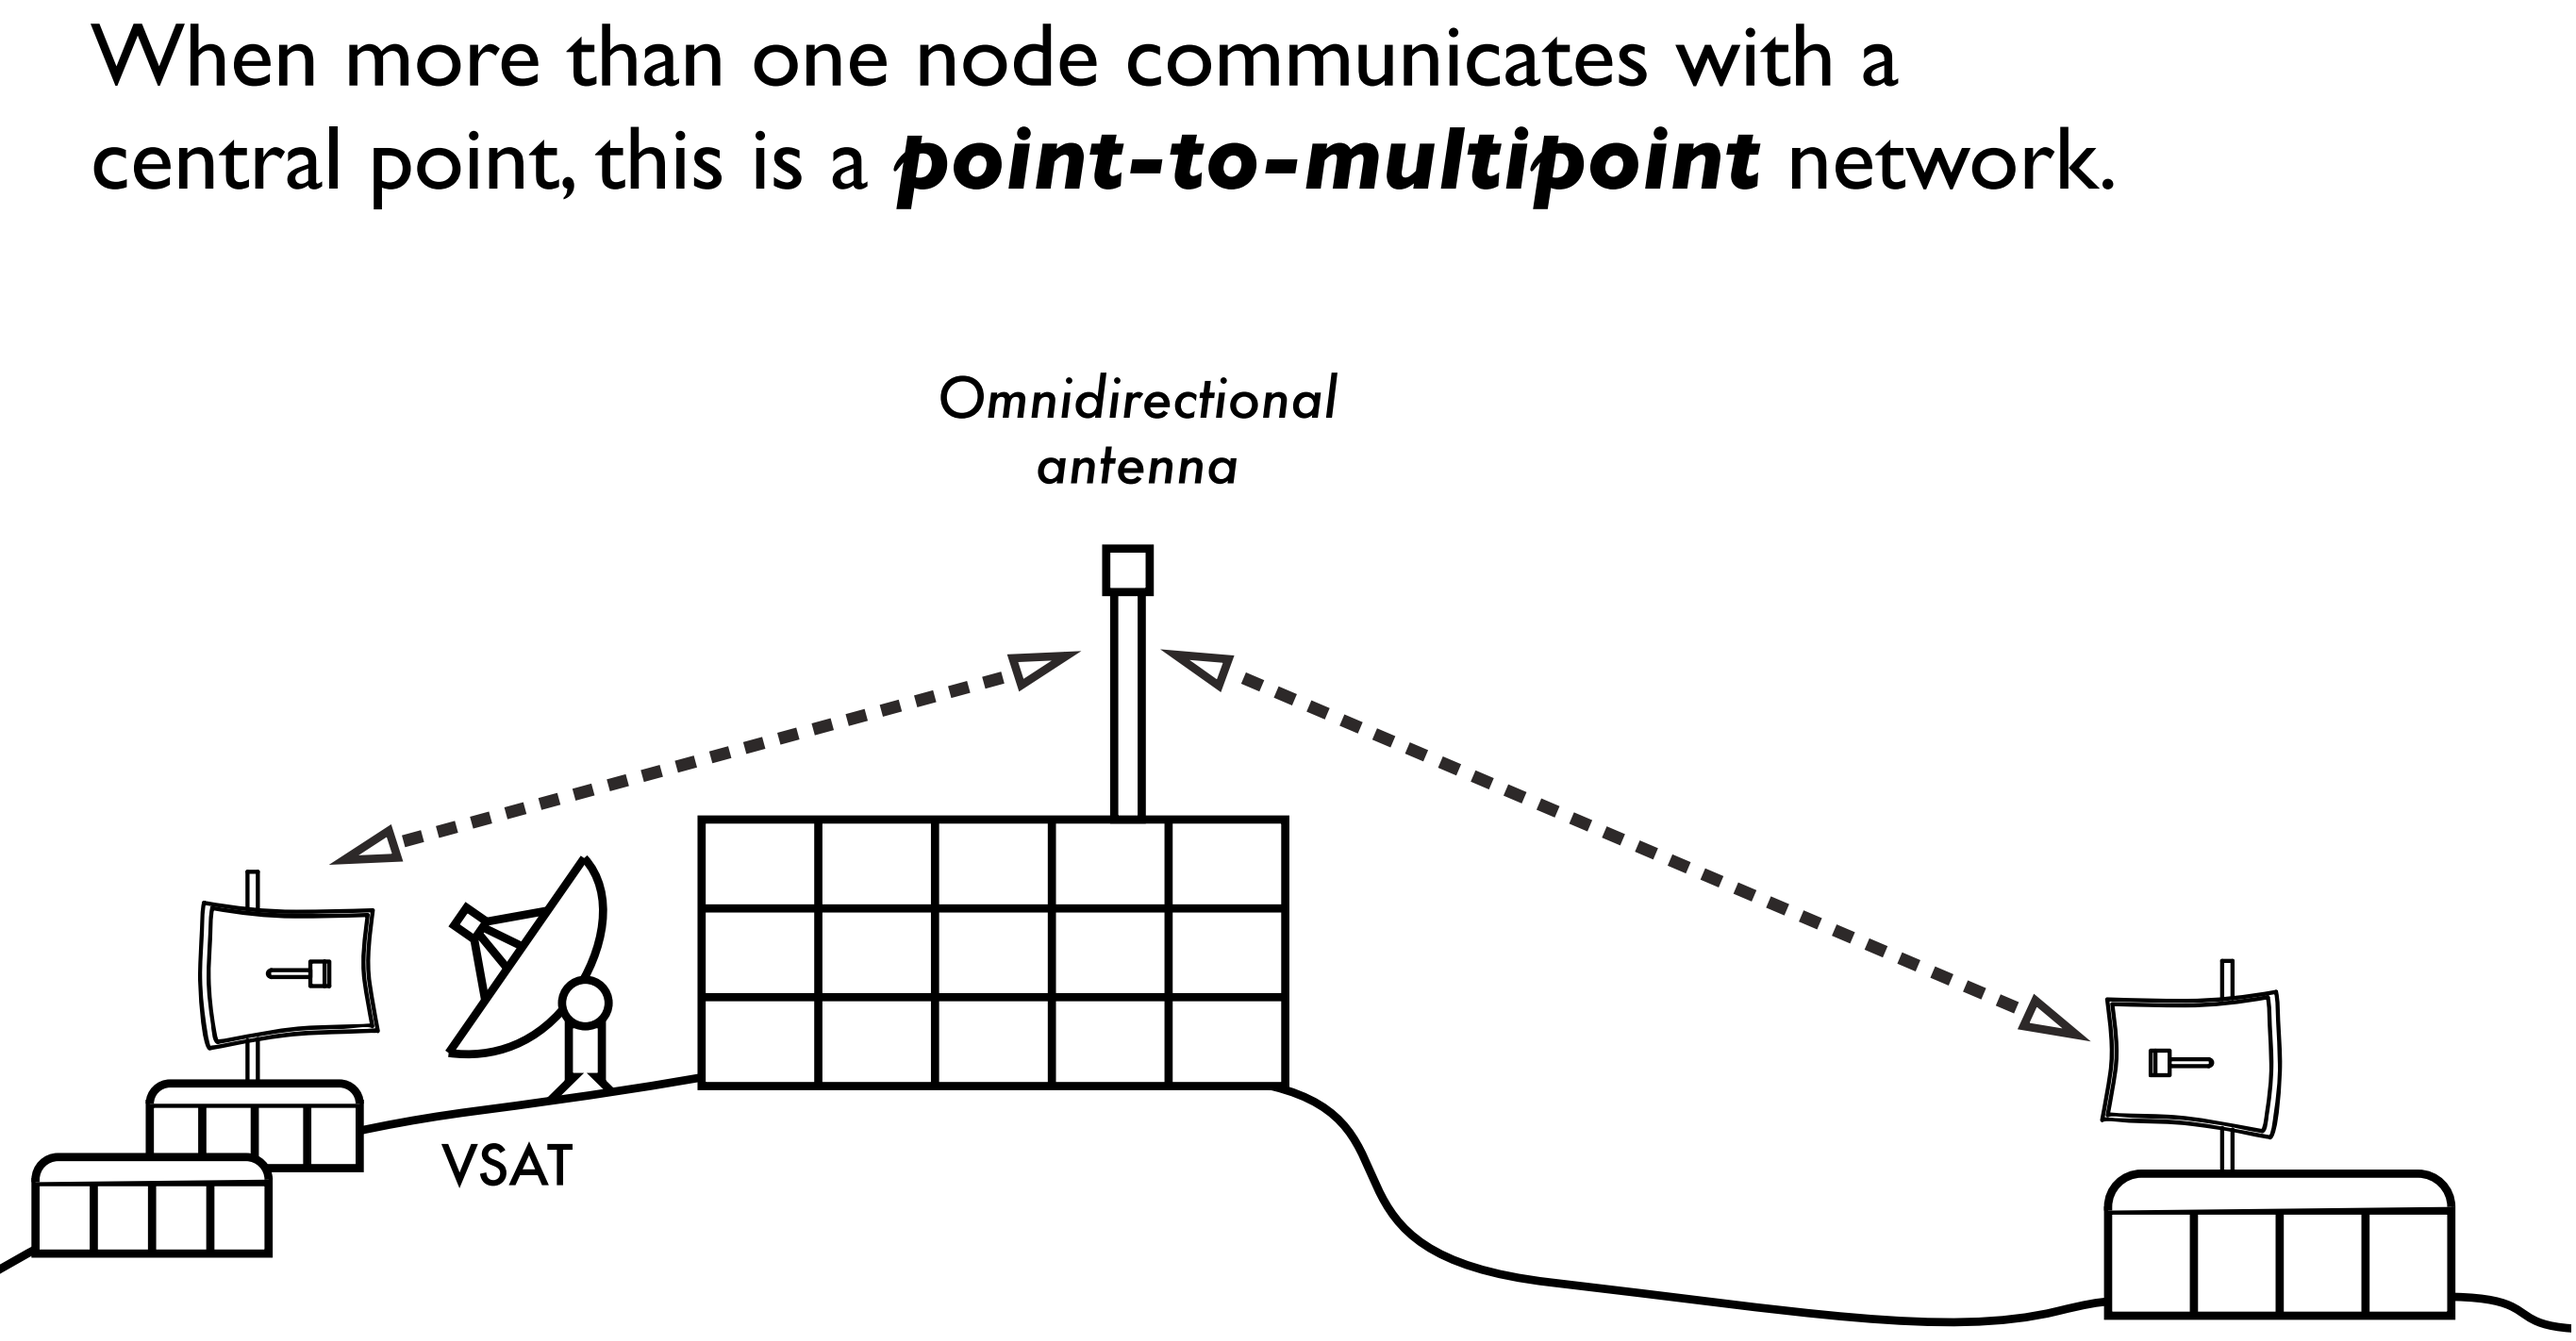

The point to multipoint network is the most common kind of topology: think of an access point with many clients.

Many times, point-to-point links will evolve into point-to-multipoint networks when word gets around that it is possible to connect to the Internet over wireless.

The engineering behind point to multipoint links is very different from point-to-point links. You can't expect to simply replace a dish with an omnidirectional antenna and expect it to work. Going from a point-to-point to a point-to-multipoint network adds complexity, since you now have multiple nodes contending for network resources. The end result is that overall throughput goes down.

Multipoint to Multipoint

When any node of a network may communicate with any other, this is a ***multipoint-to-multipoint*** network (also known as an ***ad-hoc*** or ***mesh*** network).

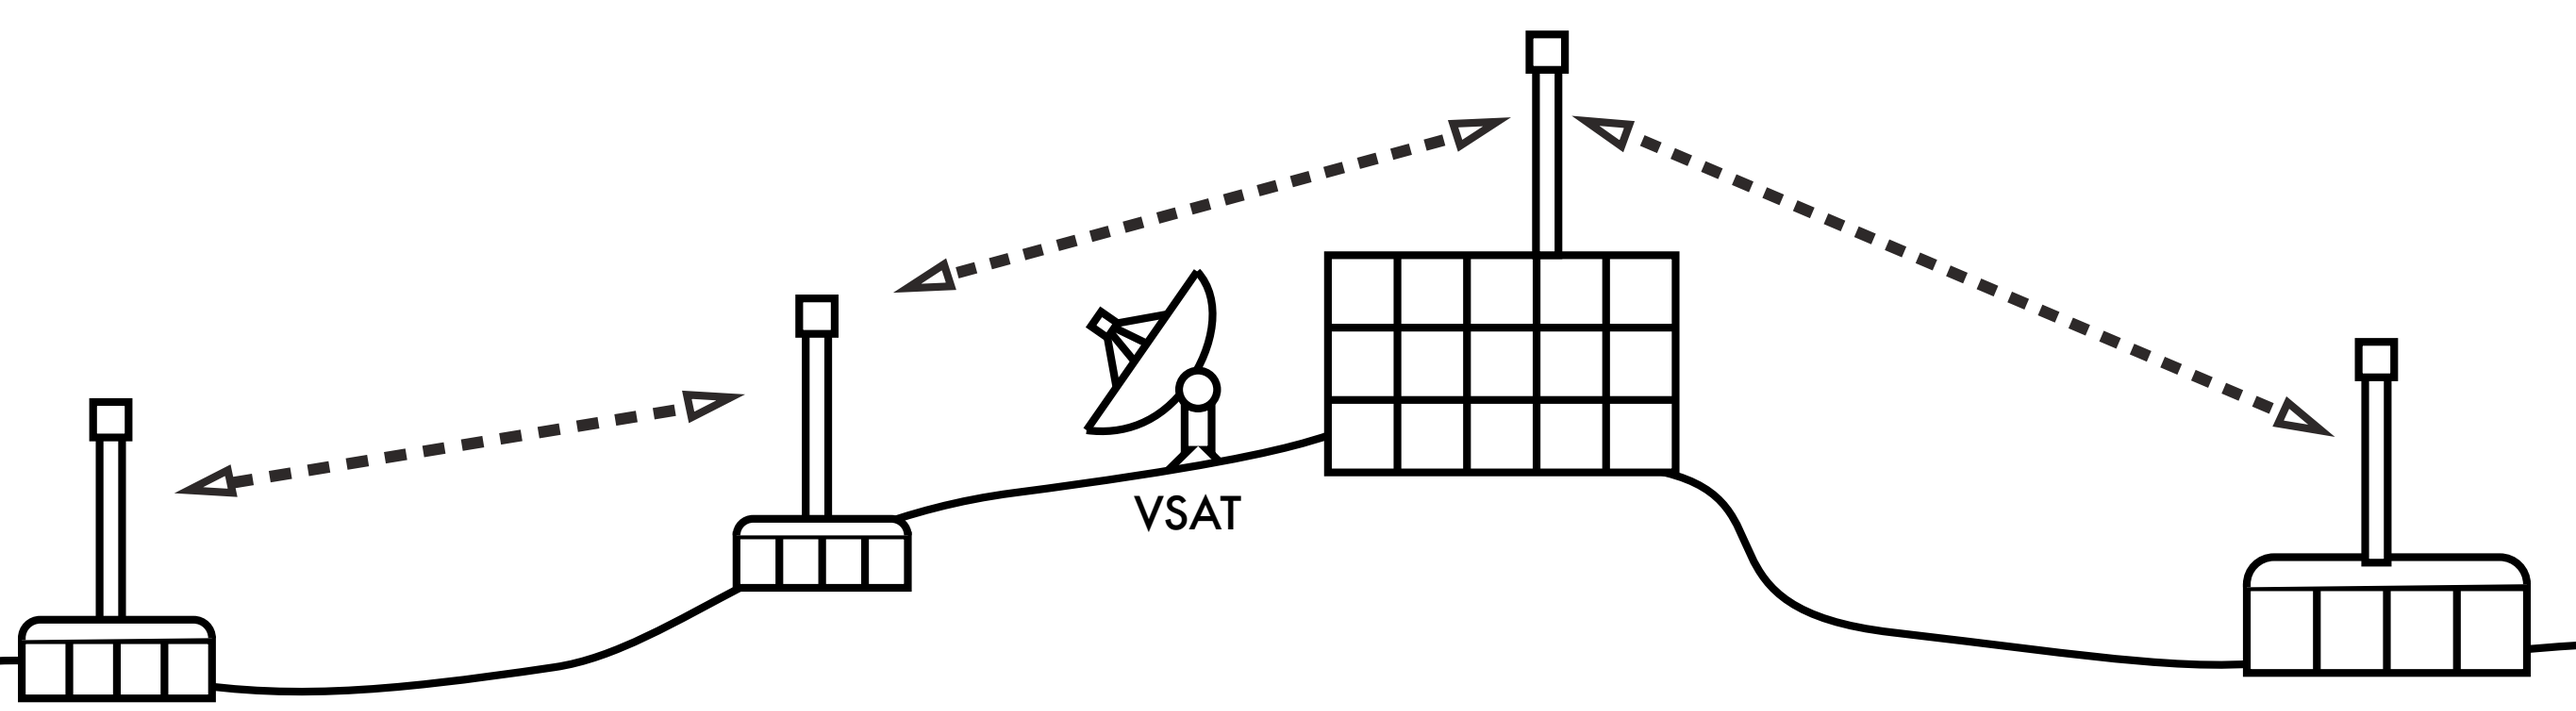

Multipoint-to-multipoint networks are considerably more complex, but also much more flexible than point-to-multipoint networks. There is no central authority in mesh networks. The mesh protocol automatically adds new nodes as they come online, without the need to change the configuration of any existing nodes.

Mesh networks can be difficult to tune compared to point-to-point and point-to-multipoint networks. One obvious difficulty is choosing a channel to use for the network. Since every node communicates with each other, only one channel can be used for a given mesh. This greatly reduces the maximum possible throughput.

WiFi radio modes

WiFi devices can be operated in one of these ***modes***:

- ***Master*** (access point)
- ***Managed*** (also known as ***client*** or ***station***)
- ***Ad-hoc*** (used for mesh networks)
- ***Monitor*** (not normally used for communications)
- Other proprietary non-802.11 modes (e.g. Mikrotik Nstreme or Ubiquiti AirMAX)

Each mode has specific operating constraints, and radios may only operate in one mode at a time.

WiFi radios can operate in one of these four modes.

Most devices can only operate in one mode at a time, so it cannot (for example) act as an access point and a client at the same time.

But you can have a wireless router that accepts more than one radio in which case you can have a radio acting as an Access Point and another as a client. This is often used to enhance the throughput of mesh networks.

In managed mode, the radio channel is not specified. It scans all channels and changes to match the AP automatically (choosing the one with the strongest signal). In ad-hoc mode, a radio is configured with a name and channel, and the radio will not change channels to match other ad-hoc radios with the same name.

Master mode

Master mode (also called AP or infrastructure mode) is used to provide an infrastructure with an access point connecting different clients. The access point creates a network with a specified name (called the ***SSID***) and channel, and offers network services on it. WiFi devices in master mode can only communicate with devices that are associated with it in ***managed*** mode.

SSID means "Service Set IDentifier". It sometimes called also ESSID (Extended SSID) or BSSID (Basic SSID), with the same meaning (not really, but the difference is not important for us now...).
For more details, see: *http://en.wikipedia.org/wiki/Service_set_(802.11_network)*

Access points create point-to-multipoint WiFi networks. A radio operating in master mode acts as an access point, advertising a network of a certain name on a certain channel, and allows wireless clients to connect to it. There may be limitations on the max number of clients allowed (this limit varies across AP models).

Managed Mode

Managed mode is sometimes also referred to as ***client mode***. Wireless devices in managed mode will join a network created by a master, and will automatically change their channel to match it.

Clients using a given access point are said to be ***associated*** with it. Managed mode radios do not communicate with each other directly, and will only communicate with an associated master (and only with one at a time).

Sometimes a device working in client (managed) mode is also called "station" or "CPE" (Customer-premises equipment or customer-provided equipment).

Clients can only talk to one AP at a time.

Ad-hoc Mode

Ad-hoc mode is used to create mesh networks with:

- No master devices (APs)
- Direct communication between neighbors

Devices must be in range of each other to communicate, and they must agree on a network name and channel.

Ad-hoc mode is used to create mesh networks. This is done by creating a multipoint-to-multipoint network when there is no master available.
Ad hoc mode can also be used to establish a link between to WiFI equipped laptops without using an Access Point.

In ad-hoc mode, each wireless card communicates directly with its neighbors.
Some vendors do not implement properly the ad hoc mode thus impairing interoperability.

Monitor Mode

Monitor mode is used to passively listen to all radio traffic on a given channel. This is useful for:

- Analyzing problems on a wireless link
- Observing spectrum usage in the local area
- Performing security maintenance tasks

Monitor mode is used by some tools (such as Kismet) to passively listen to all radio traffic on a given channel.
This is useful for analyzing problems on a wireless link or observing spectrum usage in the local area.

Monitor mode is not used for normal communications.

WiFi radio modes in action

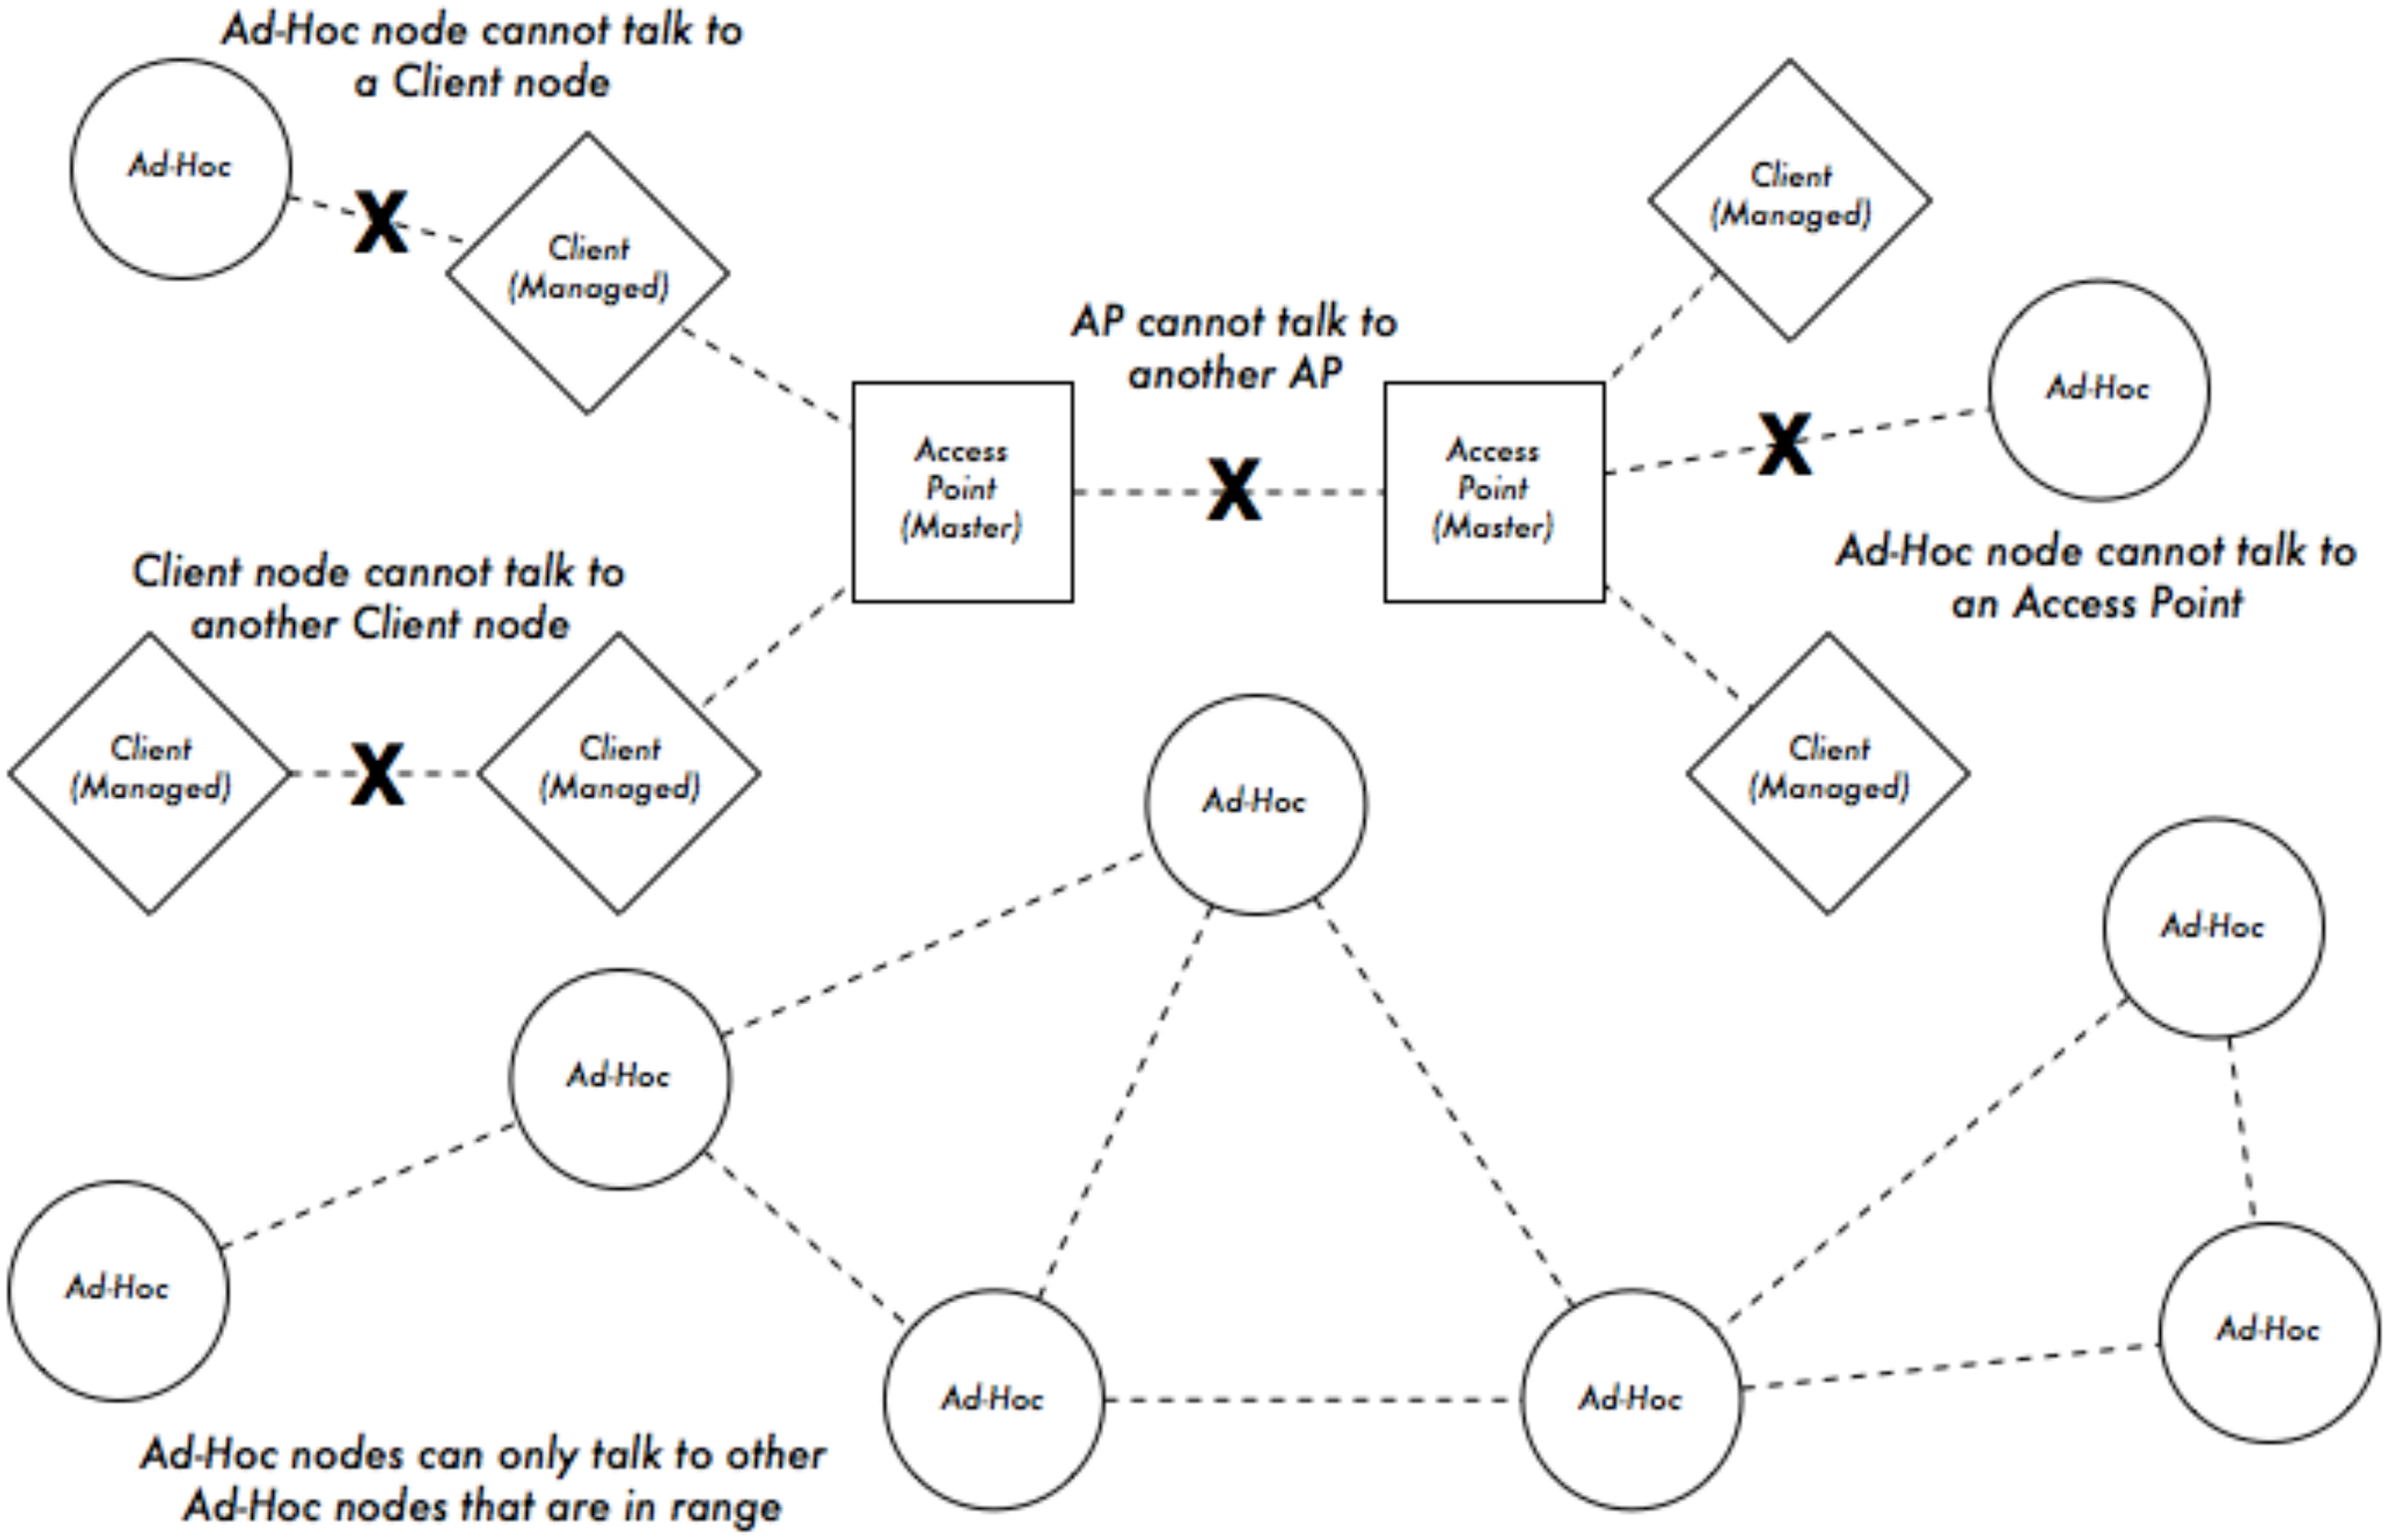

Access points cannot communicate with other access points over the air. A feature of WiFi called WDS (Wireless Distribution System) can allow up to five APs to communicate with each other via wireless, but throughput is greatly reduced, and vendor interoperability problems are common, so relying on it is generally discouraged.

Clients cannot communicate with other clients, unless they are in range of a common AP. It is a common problem that two laptops may be sitting in the same room but one can get on the network while the other cannot. If the AP is far away, and one laptop has a better antenna than the other, then only the better laptop may be able to access the network. Even though the clients are in range of each other, this does not help. Clients must be in range of an access point to use the network.

Ad-hoc nodes can only talk to other ad-hoc nodes that are in range.

Wireless Distribution System (WDS)

It is possible to allow Access Points to communicate with each other directly by using the WDS protocol. It can be useful, but it has several limitations.

- WDS may not be compatible with equipment from different vendors.
- Since WiFi is half-duplex, the maximum throughput is halved at each hop.
- WDS only supports a small number of connected APs (typically five).
- WDS cannot support some security features, such as WPA encryption.

WDS is mentioned in the WiFi standard, but the actual implementation is not fully specified, therefore different vendors often use incompatible implementations.

Routing traffic

802.11 WiFi provides a link-local connection. It does ***not*** provide any routing functionality! Routing is implemented by higher level protocols.

TCP/IP Protocol Stack		
5	Application	
4	Transport	
3	Internet	
2	Data Link	} WiFi
1	Physical	

Complex networks use some kind of routing protocol to relay traffic between nodes. WiFi only provides a link-local network connection, up through TCP/IP layer two.

Bridged networking

For a simple local area wireless network, a bridged architecture is usually adequate.

Advantages

- Very simple configuration
- Roaming works very well

Disadvantages

- Increasingly inefficient as nodes are added
- All broadcast traffic is repeated
- Virtually unusable on very large wide-area networks

The simplest network topology from layer two up is the bridge. If you bridge the Ethernet interface to the wireless, it creates a wireless “hub”, nearly the same as if every client connected an Ethernet cable to the same hub.

While the configuration is simple, it comes at a cost of efficiency, since every device on the network now shares the same broadcast domain.

Bridged access points

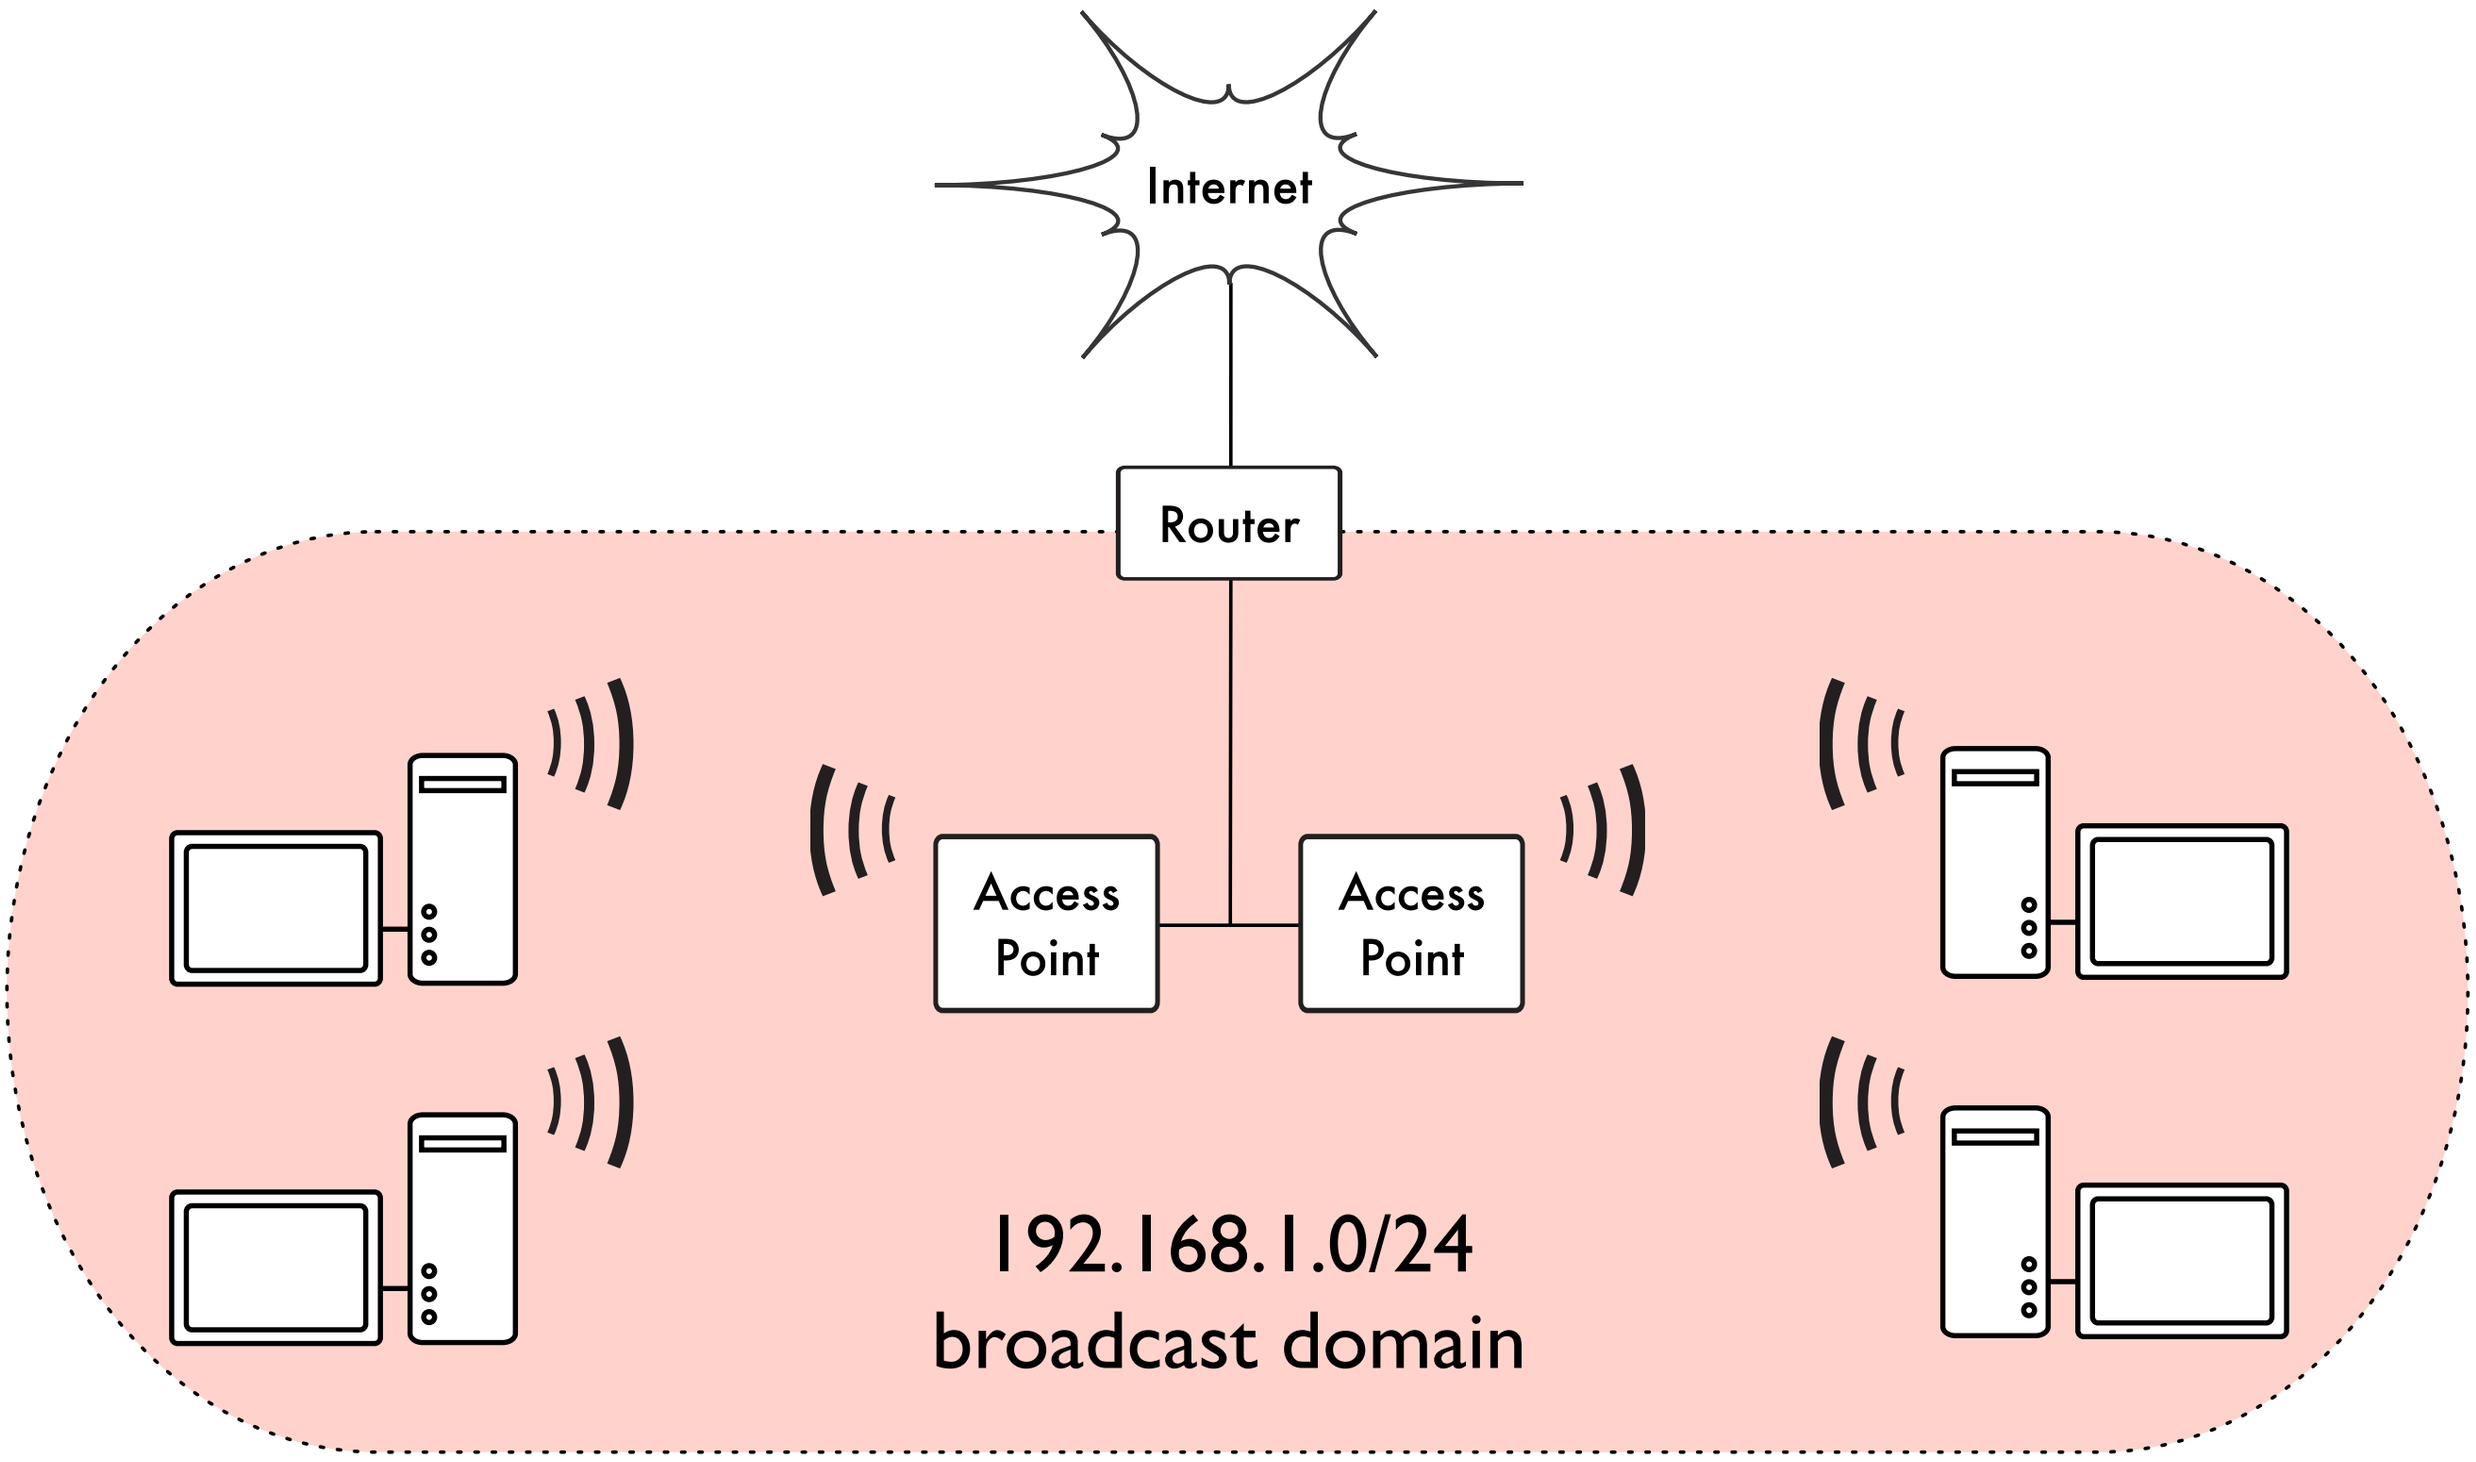

All access points in a bridged network share the same broadcast domain. All broadcast traffic (DHCP requests, ARP traffic) is sent to every node on the network. This ties up radio resources with unnecessary traffic.

Routed networking

Large networks are built by applying ***routing*** between nodes.

- ***Static routing*** is often used on point-to-point links.
- ***Dynamic routing*** (such as RIP or OSPF) can be used on larger networks, although they are not designed to work with imperfect wireless links.
- ***Mesh routing protocols*** work very well with wireless networks, particularly when using radios in ad-hoc mode.

Instead of bridging access points directly to the Ethernet, we can limit the broadcast domains to individual APs.

Traditional dynamic routing protocols work well, as long as the wireless connections are strong. Historically, a problem with slow connectivity on a link is regarded as congestion, so the protocol "backs off" and transmits less frequently to help reduce the problem. But on a wireless link, slow connectivity may be a symptom of a weak signal or interference, which is indistinguishable from congestion. Reducing the transmission rate may actually make the problem worse, while immediate retransmission may be a better strategy.

Modern mesh routing protocols (such as OLSR or BATMAN) can make use of information about link quality to make better decisions about routing and retransmissions.

Routed networking

As the network grows, it becomes necessary to use some sort of routing scheme to maintain traffic efficiency.

Advantages

- Broadcast domains are limited, making more efficient use of radio bandwidth
- Arbitrarily large networks can be made
- A variety of routing protocols and bandwidth management tools are available

Disadvantages

- More complex configuration
- Roaming between APs is not supported

Routed access points

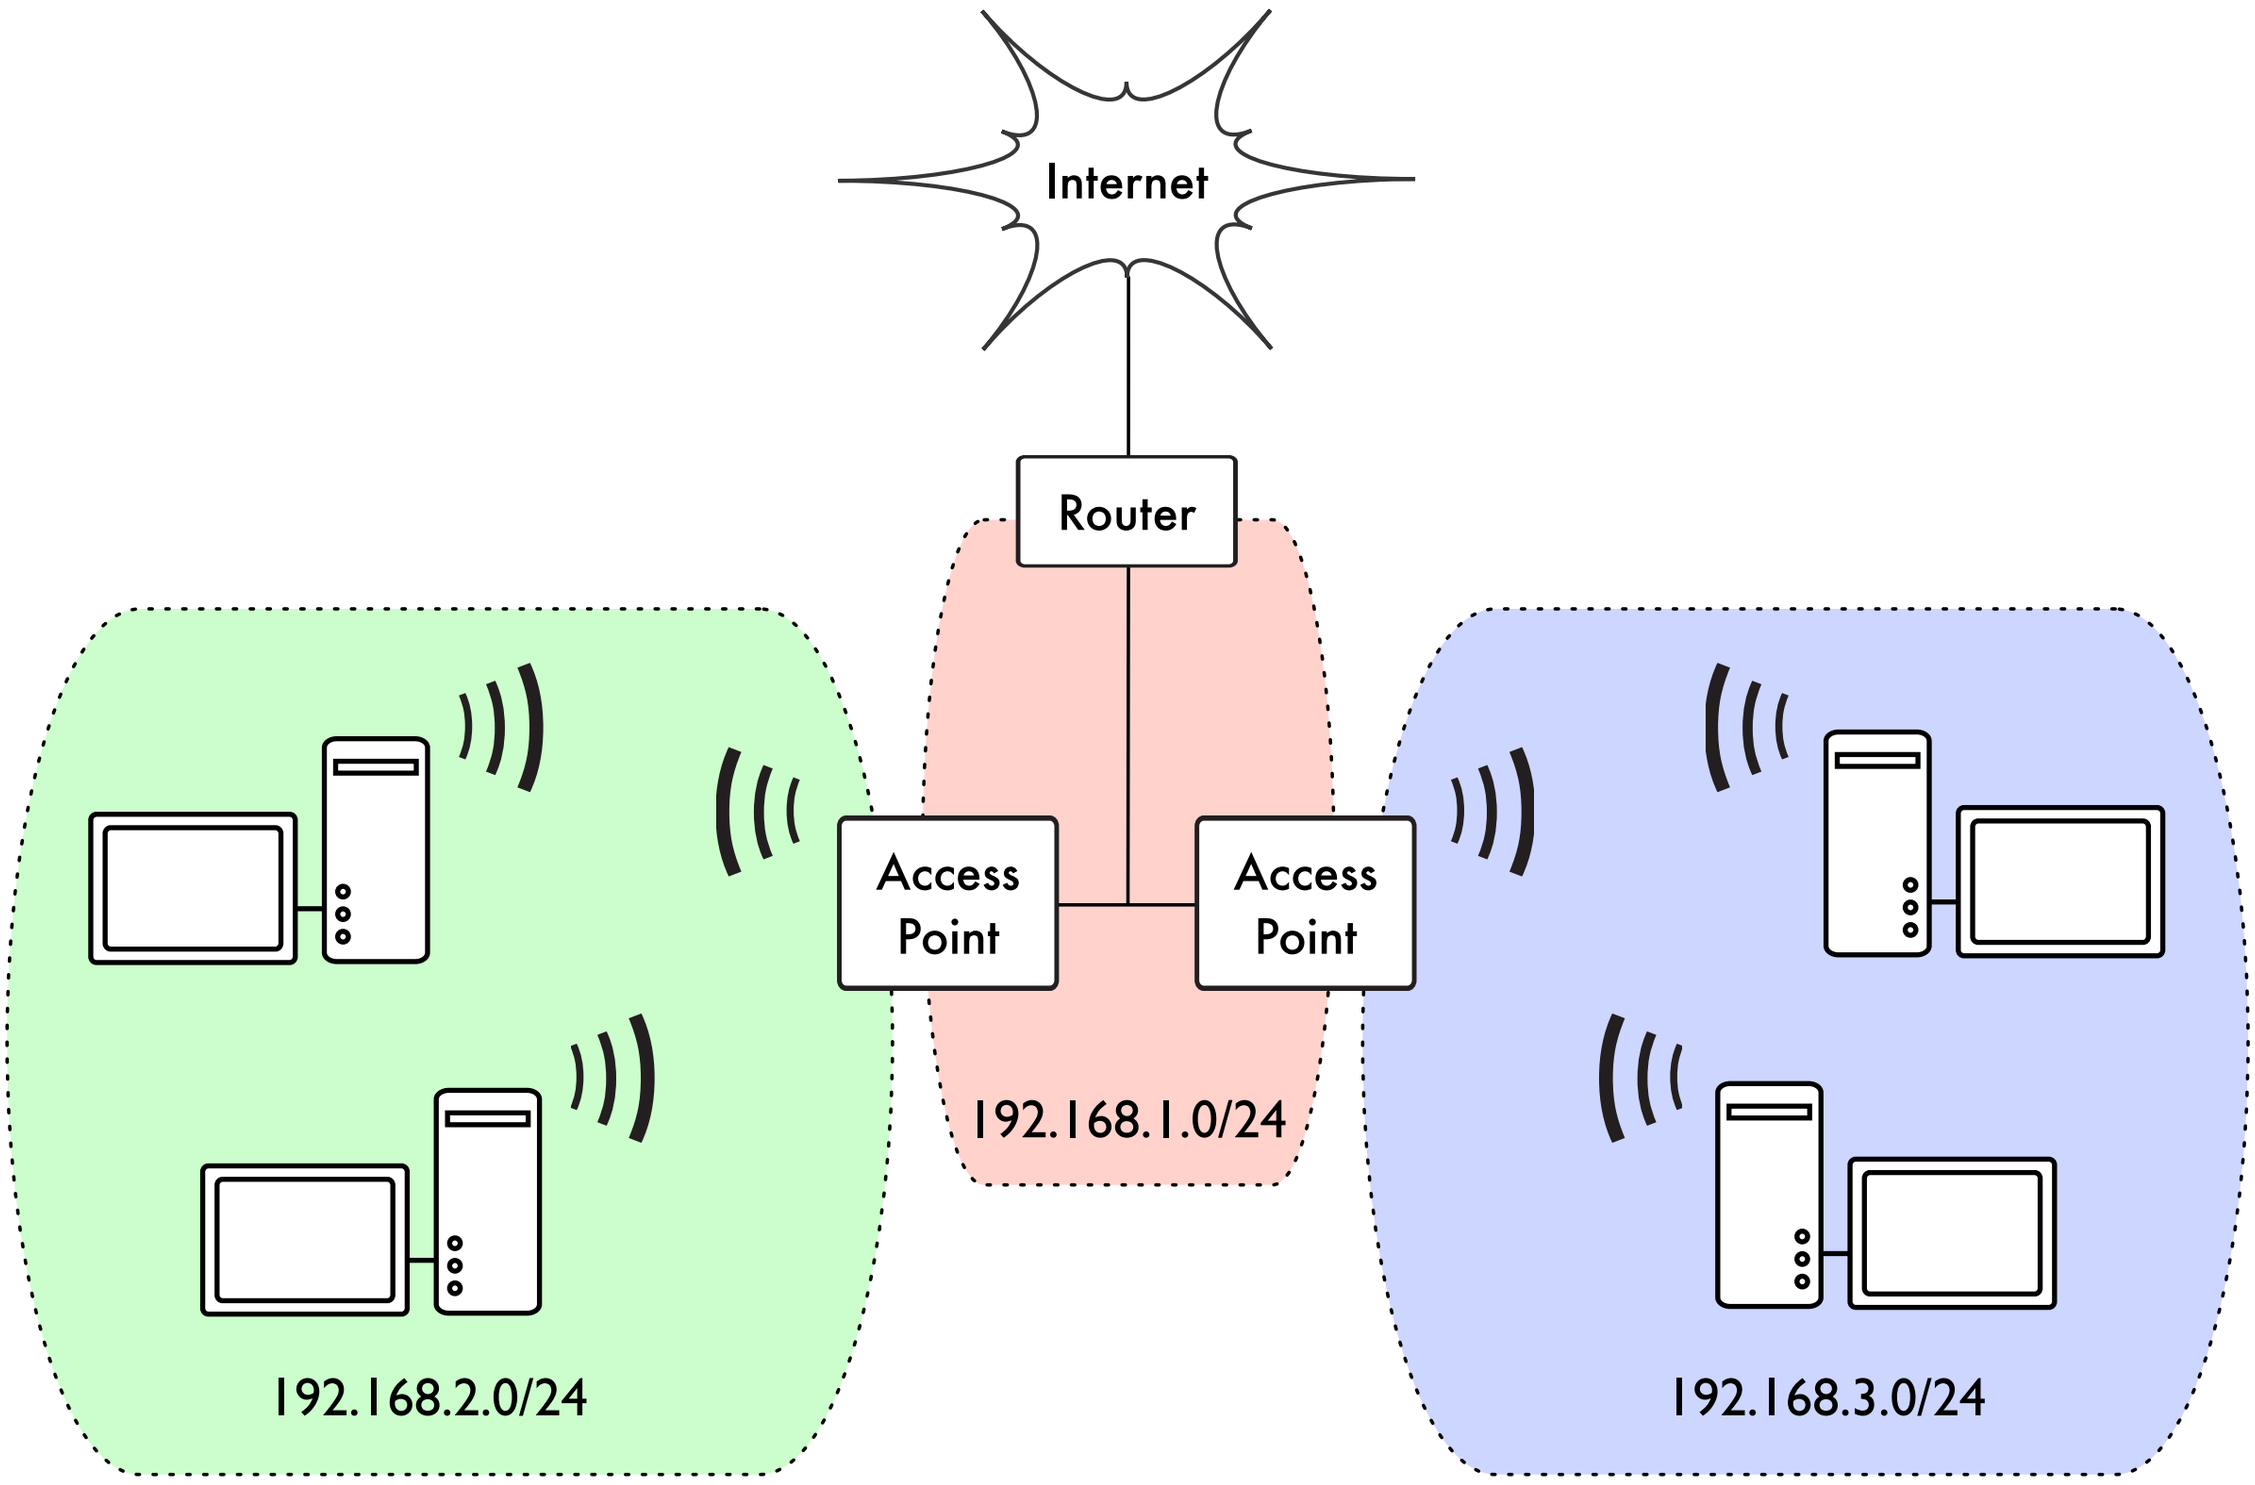

The same network can be made much more efficient by using routing instead of bridging. This reduces the size of the broadcast domains to include a single access point.

Using routing does break roaming, which is generally only a problem for IP phones or other devices that expect to maintain TCP connections when physically moving around the network.

Frequently Asked Questions

This last part of the lecture is about some common questions regarding WiFi.

Frequently Asked Questions

- How fast? (What does 54Mbps mean ???)
- How far can a network go? (the distance problem)
- How many clients can I connect to an AP?
- Are all my devices compatible?
- There are sometimes huge differences in price of APs, what should I buy?

I cannot give answers to these questions, but I hope that after this training you will be able to find the solutions by yourself. In fact, there are no "always true" solutions to these questions, the actual solutions will depend from many factors (your needs and requirements, the place you live and the availability of hardware in the local market, and many others). So you are the only one that can find out the "right" answers.

Why is it so difficult? Let analyze these questions one by one (with some humor :)

- Q: *How fast is my network? (What does 54Mbps mean ???)*
  A: The raw **radio bitrate** value (i.e. 54Mbps) specified by the standard has little to do with the actual (average) data transfer rate, because many factors may affect the performance of the network. Even worse, it's always rather difficult to perform a precise measurement of the throughput of a simple wireless link under realistic circumstances. There are tools to perform such measurements (i.e. *iperf*) but you should use them with a good knowledge of internal TCP/IP mechanisms, and analyze their results with great care.

- Q: *How far can my wireless network go? (the distance problem)*
  A: This is the "mother of all questions" in wireless networking... Simply said, the only answer that is always valid is "*from almost zero to almost infinite*" (i.e. there is no answer). You may open a discussion with all students, present a case study and analyze how a change of the following elements will affect the maximum distance of a radio link: type/gain/alignment of the two antennas, TX power, RX sensitivity (at different speeds), obstacles (fixed and variable ones), interferences, minimum required throughput, signal attenuation due to weather conditions, etc... (many other parameters can be considered, ask the students to propose a few more...). A nice exercise may be to describe a certain setup (give some hardware specs, etc.) as example, then ask to some of the students to think briefly about it and give a rough value of the maximum distance that can be achieved (call it D), then ask to other students to present a situation in which the distance can be 2D (they should imagine favorable conditions and use "good" values of parameters that were left unspecified), and finally ask to another student to imagine a situation in which the max distance is just D/2.

- Q: *How many clients can I connect to an AP?*
  A: This number depends on: characteristics of the AP, minimum accepted throughput for every client, interferences, distance of client stations, etc...

- Q: *Are all my devices compatible?*
  A: Basically, they should be all inter-compatible, as long as: they are all certified for WiFi, you use only the basic features that are defined by the standard (i.e. you forget about extra goodies like "turbo" modes, "special security" features, WDS, "boost-my-AP" settings, etc...), you configure them properly. It is a good question, indeed ;-)

A few important concepts

I can give you answers to some questions, indeed :-)

- What is a *device*?
- What is an Access Point (AP)? Can it be also a client? Are they *different* hardware?
- What is firmware? Why may I want to change it?
- I don't understand the differences between AP, device, firmware, protocols...

The answers will be given in the next slides.

A few important concepts

device (hardware)	+	firmware (O.S.+software)
hardware features: ‣ power ‣ frequencies ‣ ports (radio, ethernet) ‣ connectors, etc...		firmware features: ‣ modes (AP/client/…) ‣ supported protocols ‣ allowed max power ‣ allowed frequencies ‣ settings you can do…

all of this together makes up your AP/client

We call a "device" the combination of a piece of hardware AND its Operating System (usually called **firmware**). The OS may also include other pieces of software (e.g.: a device can run a dedicated distribution of Linux with additional packages/commands).

Alternate firmware

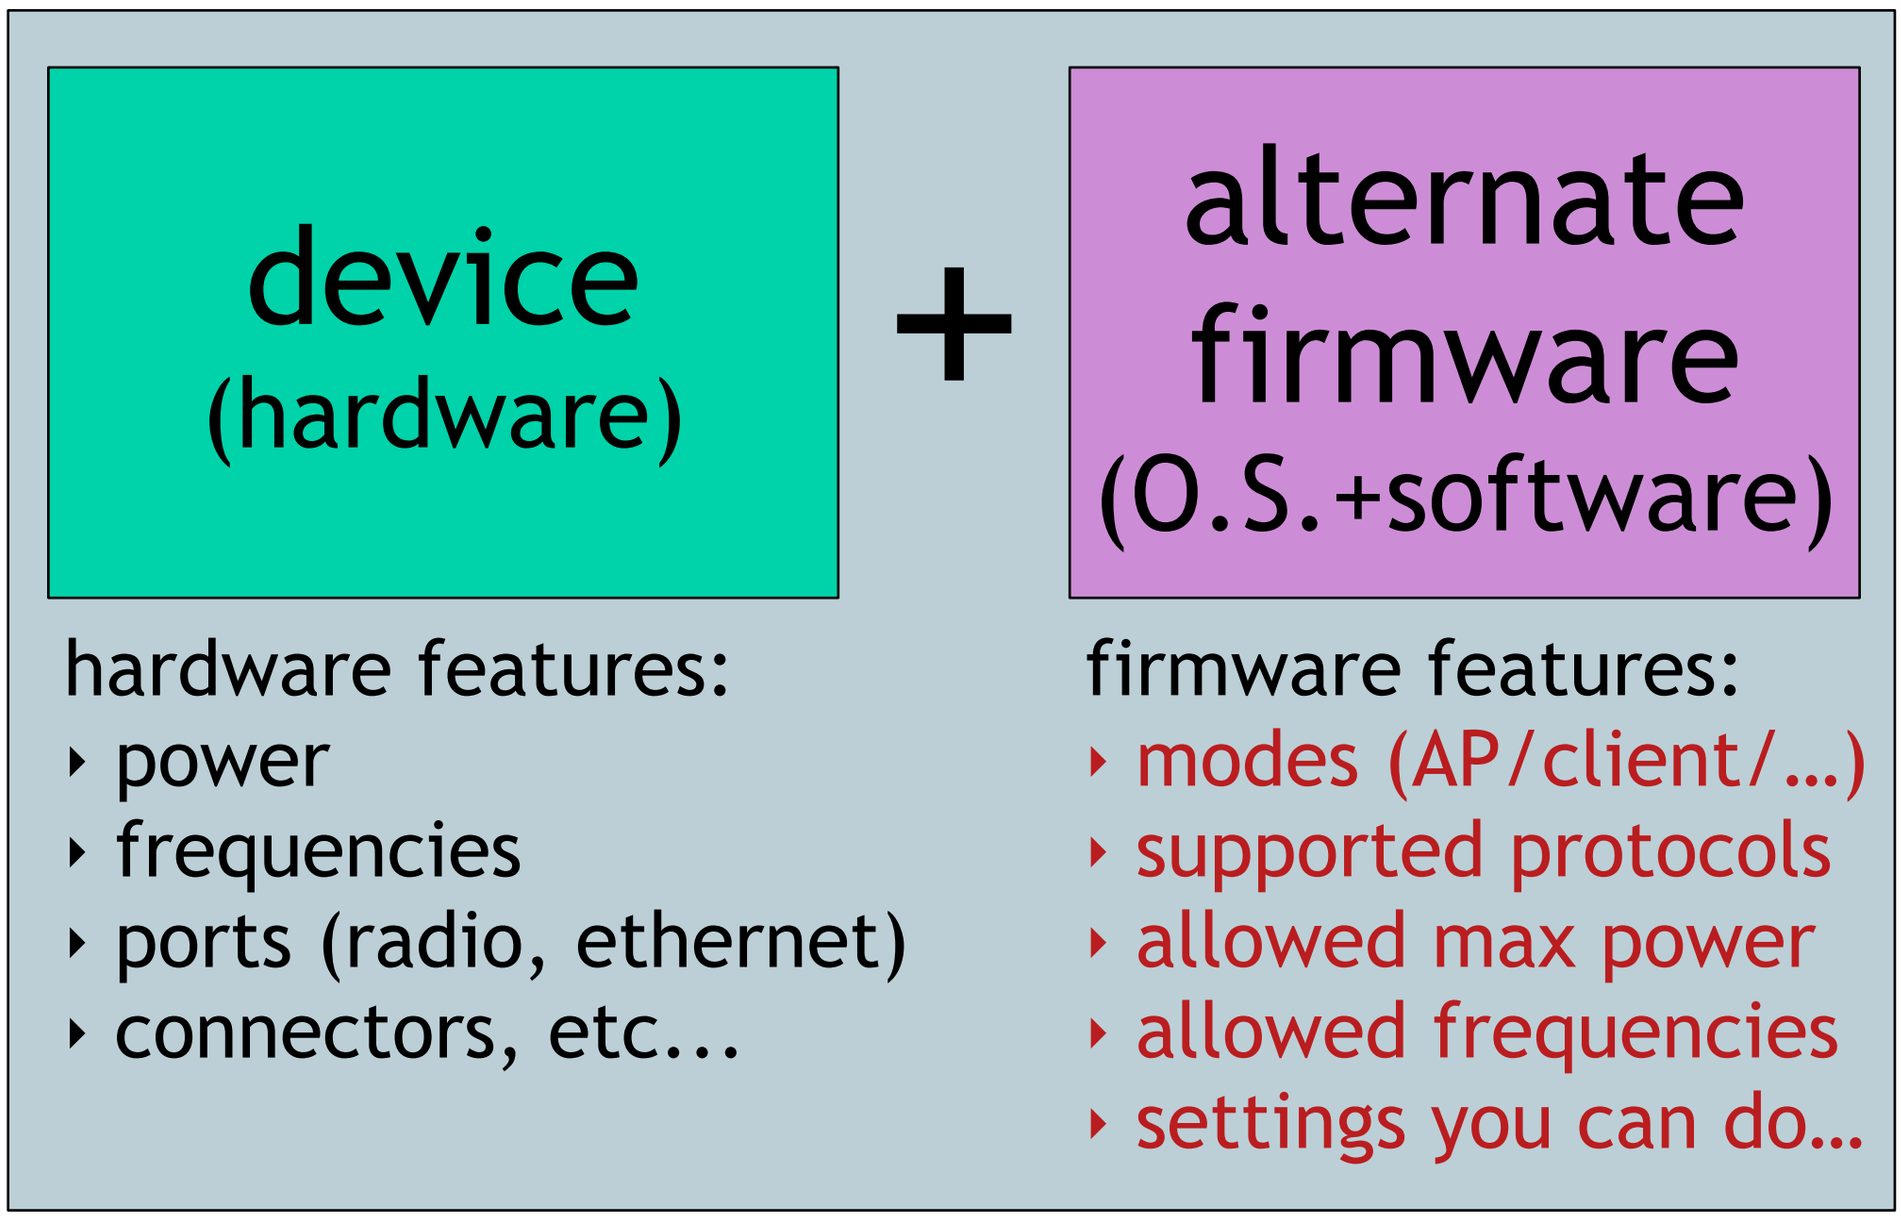

the same device *with an alternate firmware*:
it may have some *new or better features*

While we cannot change the features related to the actual hardware without changing the hardware itself, we can sometimes add or improve features by changing the firmware that run on the device.

In this example, the characteristics in red color can be different then the ones of the previous slide, because of change in the firmware.

In fact, some devices support multiple firmwares (including free / OpenSource ones), or the hardware vendor may release new improved versions of firmware for their products.

Beware of increasing the output power beyond the original specification by changing the firmware. This often is very harmful because it can cause distortion of the signal and interference in adjacent channels.

Modular hardware

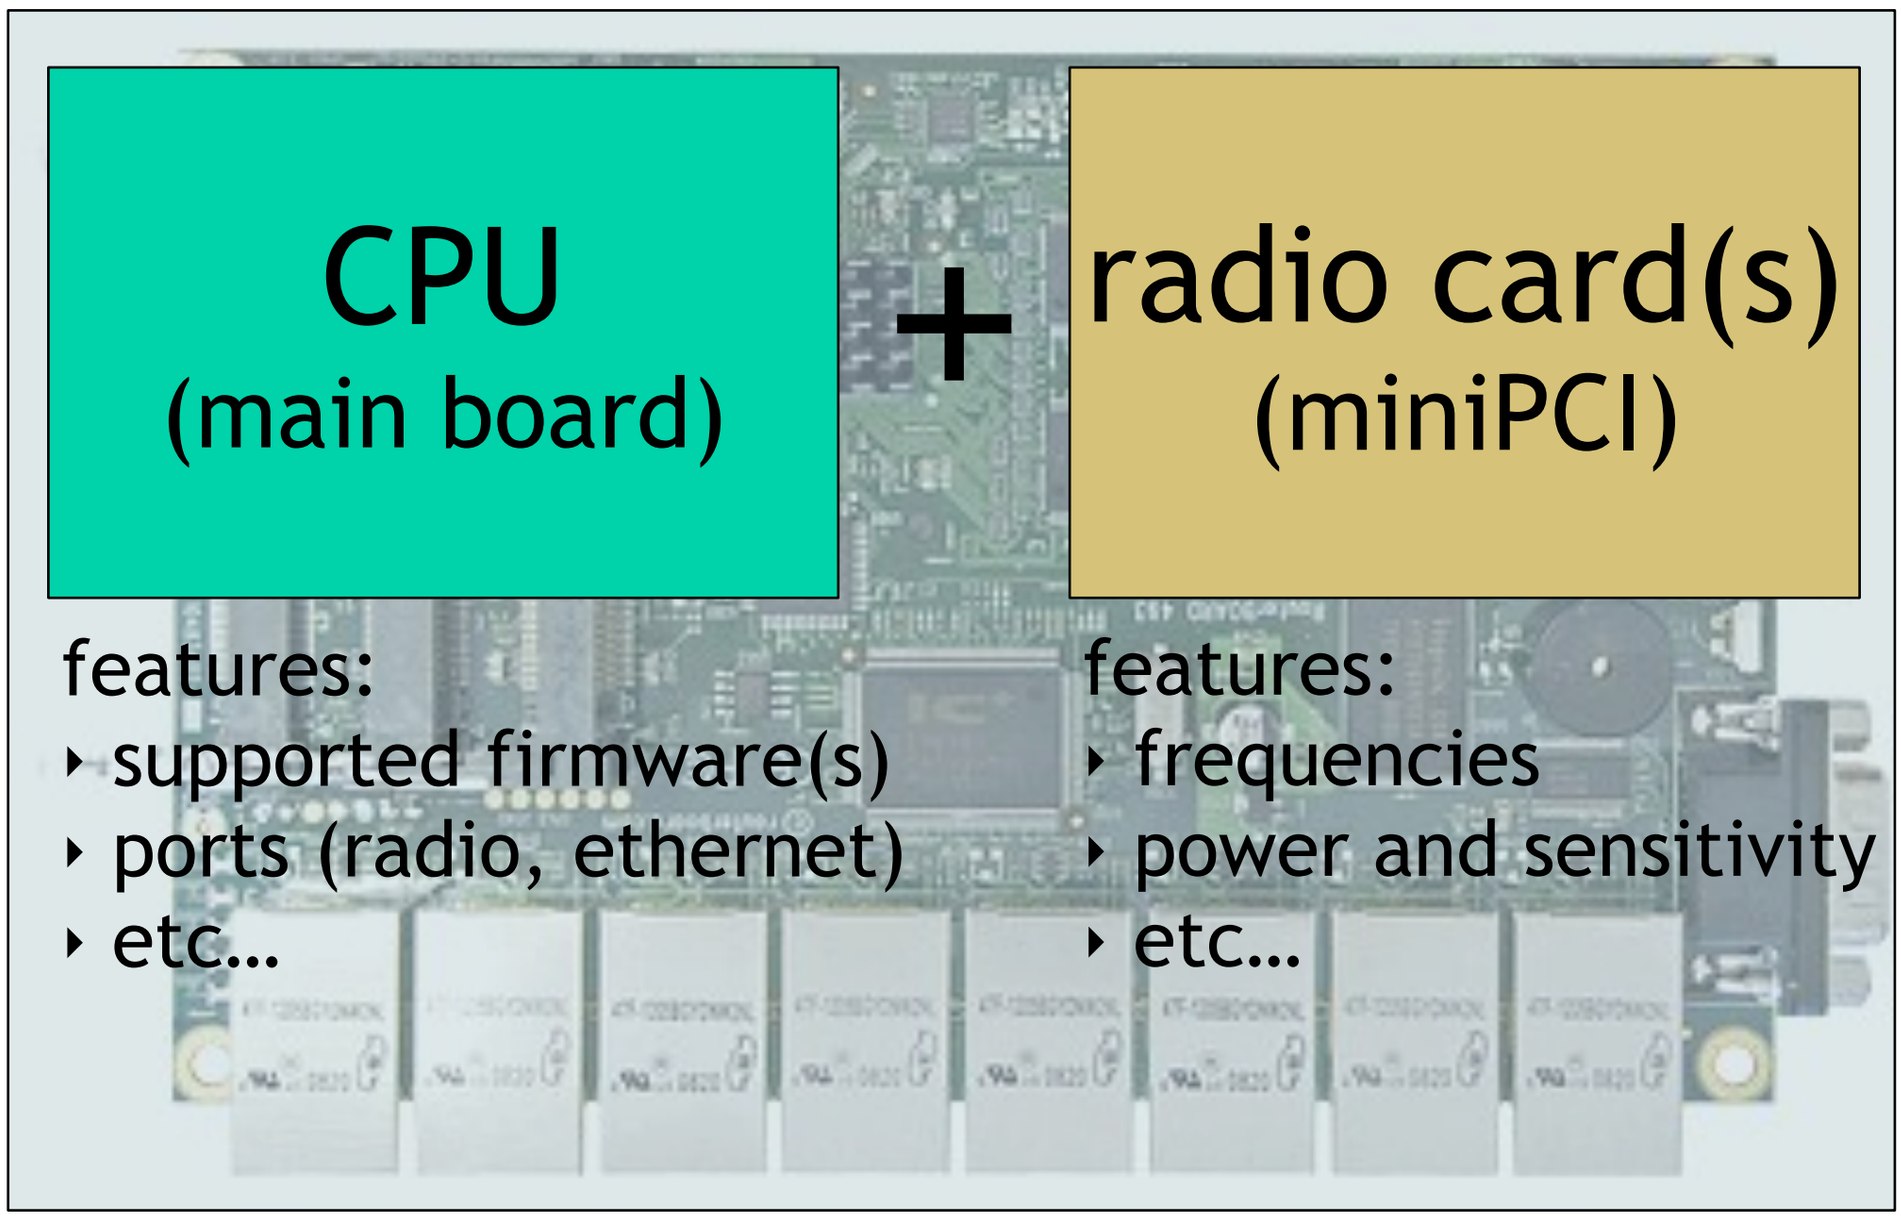

in some devices (ex: Mikrotik Routerboards)
you can change/add radio card(s)

Radio cards may come from different vendors: before buying and installing them, always check if they are compatible with the main board (and if the firmware is also compatible).

Important! You should also check if the power requirements are met (specially with cards that have high TX-power): if you install one or more high-power radio cards in a main board, you can easily overload and damage the cards, the board and/or the power supply (or the PoE injector).

A link is composed of many parts

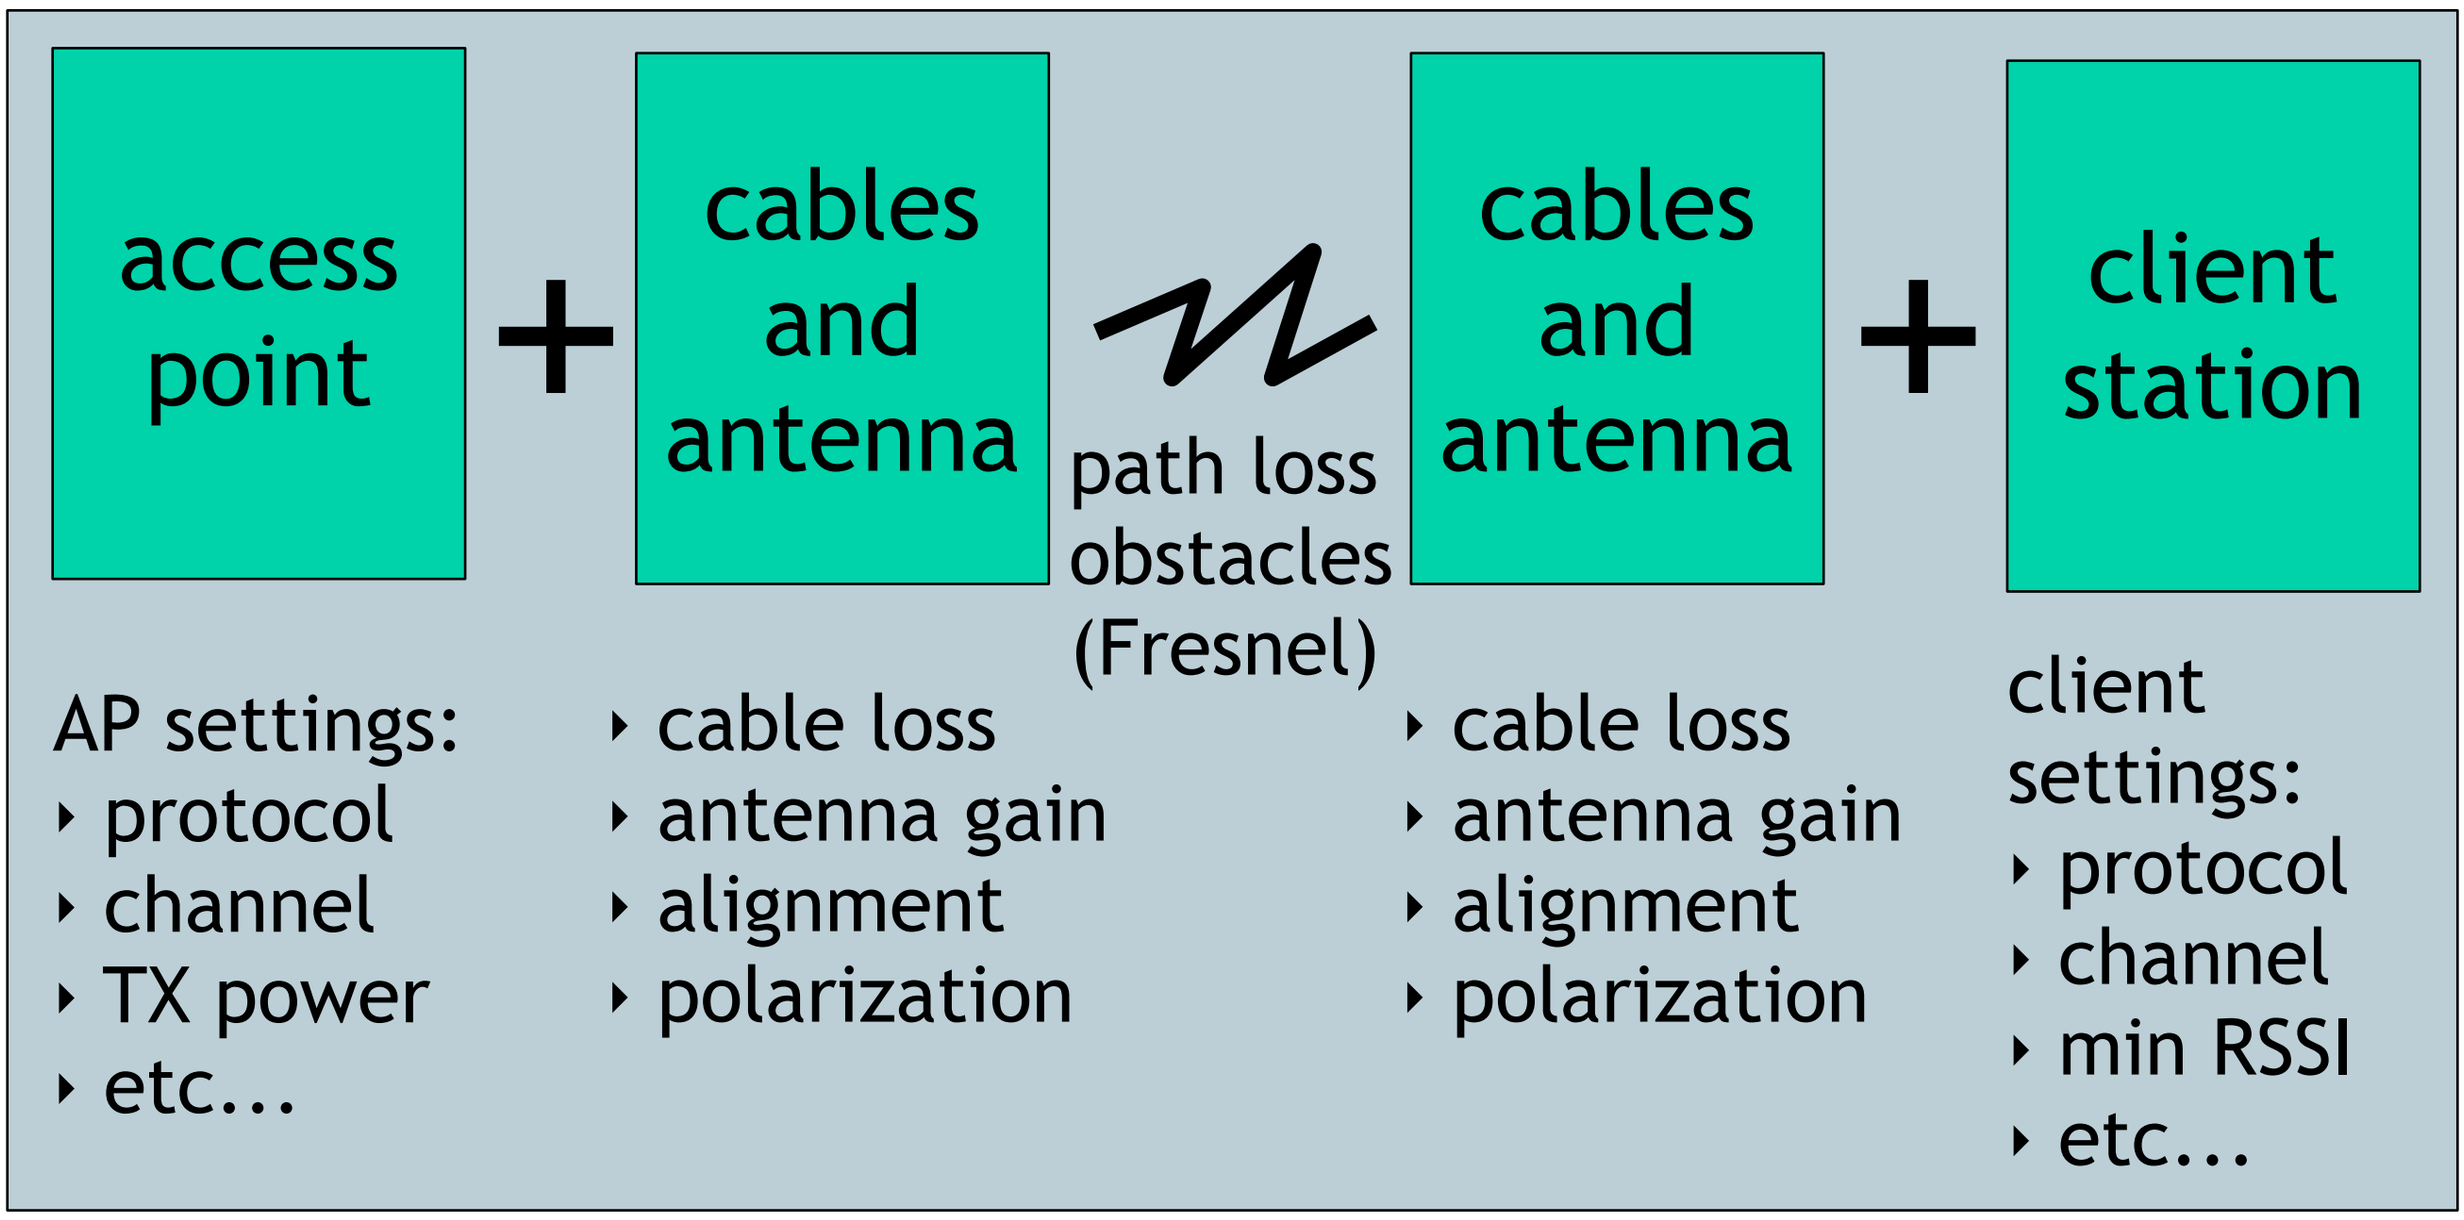

In order to have a working link: all relevant settings should match AND the link budget should allow for it

Link budget calculation will be explained in details in a separate lecture, you'll learn important concepts like path loss and Fresnel Zone.

Troubleshooting procedure should be performed level-by-level: check the TCP/IP settings, check the wireless settings, recalculate the link budget and compare the expected RX signal with the receiver sensitivity (from spec sheet), check the alignment of antennas, check all cables and connectors, check hardware.

Thank you for your attention

For more details about the topics presented in this lecture, please see the book ***Wireless Networking in the Developing World***, available as free download in many languages at:

http://wndw.net/

See Chapter 4 of the book for more detailed information about the material covered in this talk.