

Accessing Networks “Randomly”

Q2.1 What is the difference between a Basic Service Set (BSS) and Extended Service Set (ESS)?

A BSS would be in a location where there is a *single* AP, like a home (as we will see in a minute), small office, or a coffee shop. An ESS is very similar, but for larger deployments where multiple APs would be needed to provide adequate service to the users, such as a hotel or an office building. All of the APs in a single ESS will offer the same SSID.

Q2.2 What does the WiFi deployment look like in a typical home?

You can see a typical home WiFi setup in Illustration 3. Your wireless devices connect directly to a wireless router, which serves as the AP in your home’s BSS. Other devices, like desktops, can connect to your router through wired connections. Your router is then connected to a modem, which goes to the access network (*e.g.*, cable) extending outside your home.

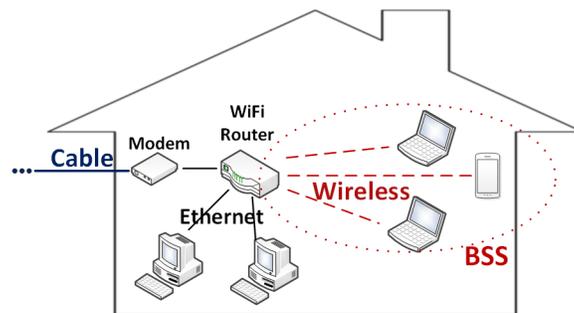


Illustration 3: Your house probably has a wireless router, which is a router with an integrated access point. Both wireless and wired devices connect to it. The wired devices, like your desktop, connect to it via Ethernet cables. The router is connected to a modem, in turn to the access network like DSL, cable, or fiber. More recently, all three devices – the modem, router, and AP – may be found within the same device.

Q2.3 Should I use cellular or WiFi for the Internet on my smartphone?

Your smartphone can offer two options to connect to the Internet: cellular or WiFi. Which to choose? Well, if there's no WiFi around, you're forced to stick with 4G (or 3G). You find this especially true when you're in transit. Cellular operators have done a good job to ensure that you can get reception pretty much anywhere you go. But when you're at a hotspot, you typically prefer WiFi. Of course, this depends on whether it is free and accessible or not.

If you're sitting next to the AP with not many people around, the WiFi data rate can be very impressive. But if you sit outside the limited range of the AP or move outside of your current BSS, you can easily get disconnected. And if you share the air with ten or so other WiFi devices, the speed can drop substantially. You've probably experienced this in a crowded public WiFi hotspot.

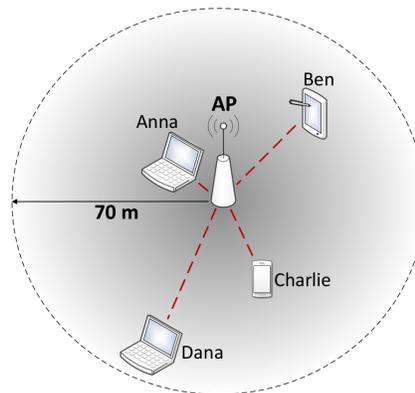


Illustration 4: Illustration of a few people in the same BSS. The coverage area (for 802.11n) extends roughly 70 m indoors (more like 250 outside). A person like Anna sitting right next to the AP will have much better signal strength than someone like Dana at the edge of the BSS. Interference also affects the channel conditions: if any two people transmit at the same time, their signals will collide.

We illustrate differing channel conditions in Illustration 4. Anna, who is right next to the AP, will have better signal quality than someone like Dana who is at the edge. And the devices cause interference for each other: think of any two trying to send to the AP at the same

time. These transmissions will collide, just like a car crash resulting from someone running a stop sign.

A real life example of this happened when Steve Jobs, the former CEO of Apple, introduced the iPhone 4 at a conference in 2010. His demo phone had trouble connecting via WiFi since there were a lot of people in the auditorium. He had to ask all the attendants to get off the WiFi before he was able to get on.

Q2.4 How does my device select an Access Point?

A WiFi device needs to make sure it has chosen the best access point for itself. Which one is the “right” one? Well, among all the ones that it can access, it wants to make sure it has the highest signal-to-interference ratio possible. So, the device will regularly scan the air and determine whether its current AP offers the highest SIR. If not, it will switch. We call the AP it chooses the one it *associates* with. This would occur in an ESS if the client device is roaming; in this case, the session can be handed from AP to AP with no service interruption.

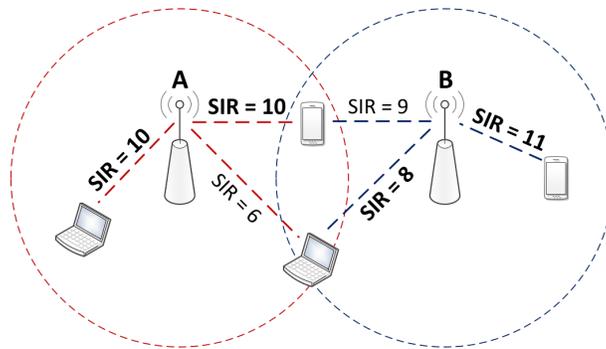


Illustration 5: Our WiFi devices regularly scan the area to determine whether they are associating with the best AP or not. The best AP is the one to which the SIR is the highest. The phone and laptop in the middle would rather be with A and B, since $10 > 9$ and $8 > 6$, respectively. The laptop on the left has only one choice, A, and likewise, the phone on the right must use B.

In Illustration 5, we give an example of this. The laptop on the left

and the phone on the right only have one choice: AP's A and B, respectively. For the devices in the middle, the cell phone will choose A ($10 > 9$), and the laptop will choose B ($8 > 6$).

Q2.5 Which channel will my device use to communicate with its AP?

Each AP operates on a given frequency channel. The overall ISM band is divided into a number of such channels, each with a pre-determined width. Taking the example of 802.11b (in the US), there are 11 located around 2.4 GHz, each with a width of 0.022 GHz (22 MHz). On each side of a channel's **center frequency**, it occupies half of this, or 0.011 GHz (11 MHz).

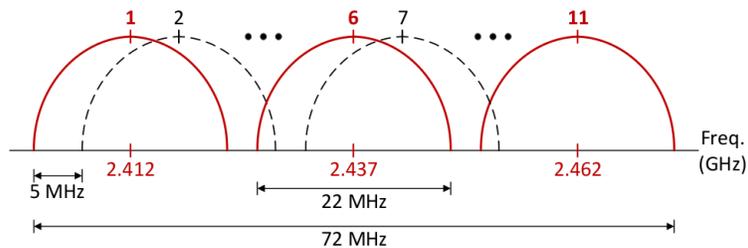


Illustration 6: Illustration of the frequency range for WiFi's 802.11b, around 2.4 GHz. Each AP has a designated channel that they transmit and receive on, which is how they communicate with devices. There are 11 channels, separated by 0.005 GHz (5 MHz) on the frequency axis. Channel 1 is centered at 2.412 GHz, channel 2 at 2.417 GHz, and so on. As we can see, channels 1, 6, and 11 are non-overlapping.

We show the allocation for 802.11b in Illustration 6. Notice that the channels *overlap* with one another. Why is that? Well, we only have 72 MHz to start with. In order to have 11 non-overlapping at 22 MHz each, we would need $11 \times 22 \text{ MHz} = 242 \text{ MHz}$ in total. To find the maximum number of non-overlapping channels possible with our range, we can divide the total width (72 MHz) by the width per channel (22 MHz) to get 3.27. Rounding this down, we see that there are at most 3 non-overlapping channels (1, 6, and 11).

In order for a device to communicate with its access point, it will select the channel that its AP is using. This means that all devices in

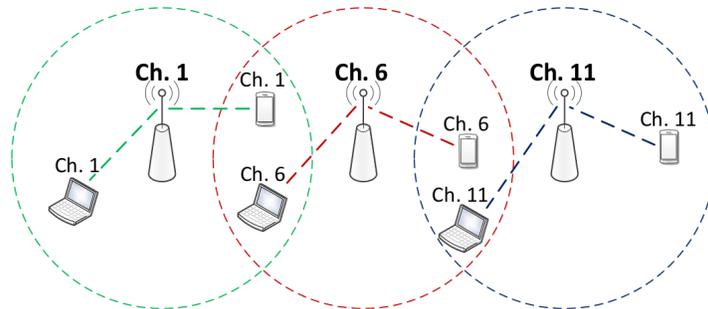


Illustration 7: Each device that wants to associate with an AP has to use its corresponding frequency channel. Here, we show three APs with coverage areas that intersect. Ideally, they would each be on non-overlapping channels, which are 1, 6, and 11 with 802.11b. This allocation would cause users in the intersection regions not to interfere. For instance, the phone and laptop between Ch. 1 and 6 are close to each other, but since they are on separate channels, they don't interfere.

the same basic service set are using the same channel. If two APs are close in proximity, then selecting the same channel causes unnecessary interference: the devices not only have to contend with others in their basic service set, but also those in the neighboring one. Ideally, we would want the access points to use different channels, as illustrated in Illustration 7. Unfortunately, many WiFi deployments just use the default settings of their access points, meaning that many are on the same channel unnecessarily, hurting performance.

Q2.6 What speed will my device achieve?

We saw that one of the distinguishing features between different versions of WiFi is the transmission rate. More precisely, this is the *maximum* transmission rate, and will only be achieved under very ideal conditions.

In many WiFi hotspots, the channel condition fluctuates as the interferers come and go. You may have 25 Mbps now, but once someone else turns their laptop on, you are down to maybe 15 Mbps. The maximum rate is rarely achieved, and the AP will tell the devices to

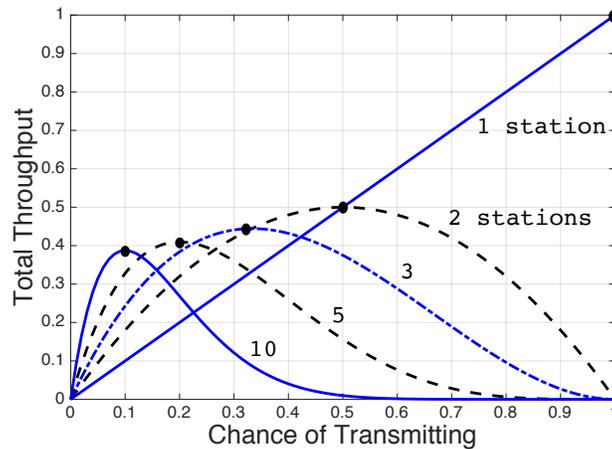


Illustration 8: Plot of ALOHA's total throughput for different numbers of stations (1, 2, 3, 5, and 10). Notice that the maximum achievable throughput drops when there are more stations. This shows how ALOHA (and random access in general) is not scalable to a large number of users.

back off to a lower rate if there is too much interference to handle the current one. The rate could be as low as 1 Mbps in 802.11b.

Q2.7 How does the ALOHA throughput change as we add more devices?

In Illustration 8, the total throughput is plotted for various numbers of stations. For one station, the total throughput keeps increasing with the transmission probability, because there's no chance of it colliding. For two stations, we start to see the effect of collision overpowering additional transmission once the probability is higher than 50%. To maximize throughput here, we should have each of the two stations transmitting 50% of the time. As we add more and more stations, we see that collisions begin to overpower the throughput for lower and lower transmission probabilities: for three stations, the maximum is reached at 33.33%, for five, it is reached at 20%.

Perhaps the most important takeaway from this graph, though, is not how the number of stations affects the optimal transmission probability, but rather how it affects the maximum achievable throughput. How does the highest point on the plot change with more stations?

For one station, it is 100%, for two it is 50%, for three it is 33.33%, and so on. The achievable throughput keeps dropping. And this is not the per-station throughput. It is the total throughput that the whole network is achieving! Adding another station makes the whole system less efficient.