Machine Learning for Networking

Lab 1

Marc Carrascosa Zamacois: marc.carrascosa@upf.edu

Lab structure

- 2 labs with 3 sessions each
- One report for each lab (after 3rd and 6th sessions)
    - Delivery date Lab 1 report **4/6/2021**
    - Delivery date Lab 2 report **31/6/2021**
- 1st and 2nd sessions:
    - 30-45 minutes of theory per session
    - 60-45 minutes of work per session
- 3rd session is only to solve doubts / questions

Labs count for 40% of the grade (20% for each lab/report)

The association problem

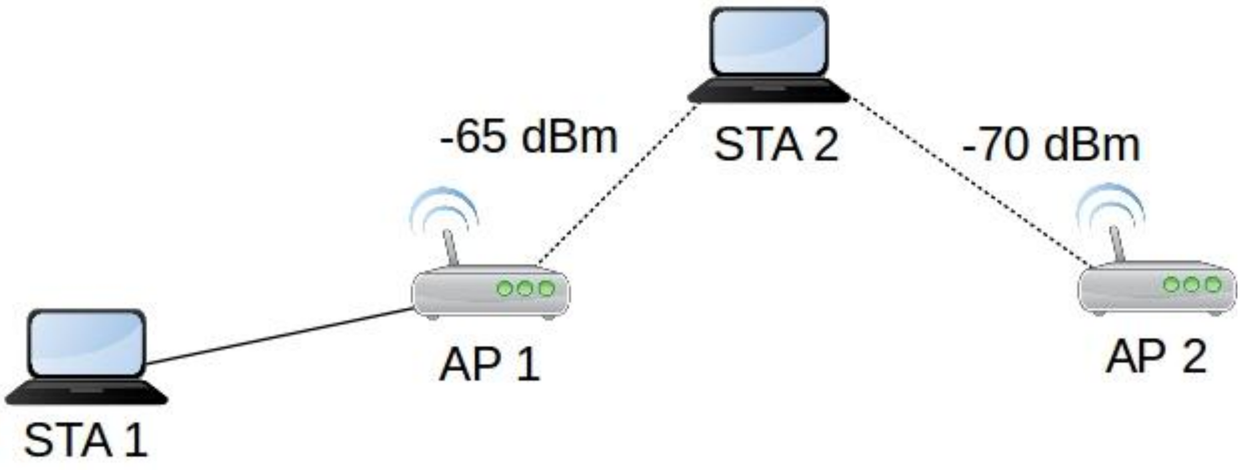

Which AP should STA 2 choose?

- AP 1 has the strongest signal
- AP 2 is the optimal, as AP 1 already has a STA

Two cases

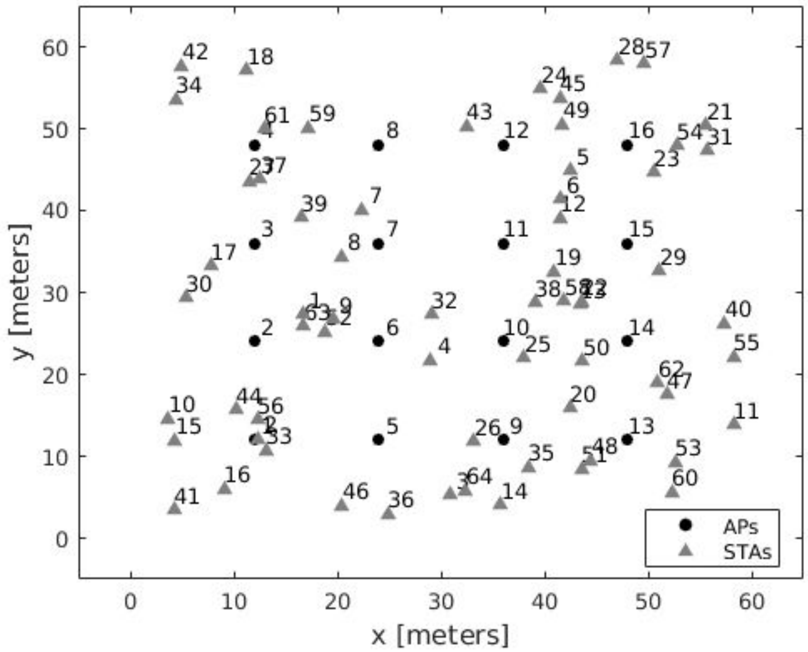

Uniform distribution of STAs

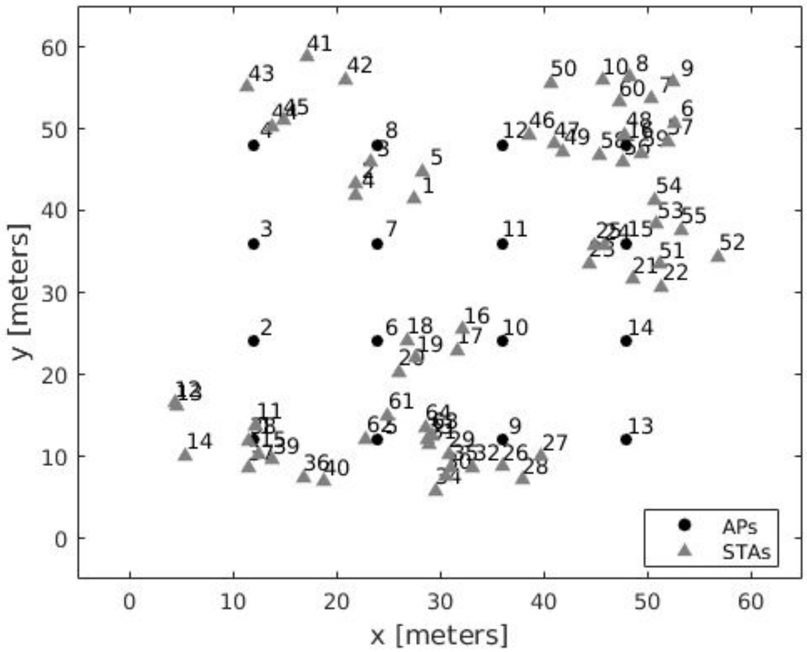

Clustered distribution of STAs

Clusters

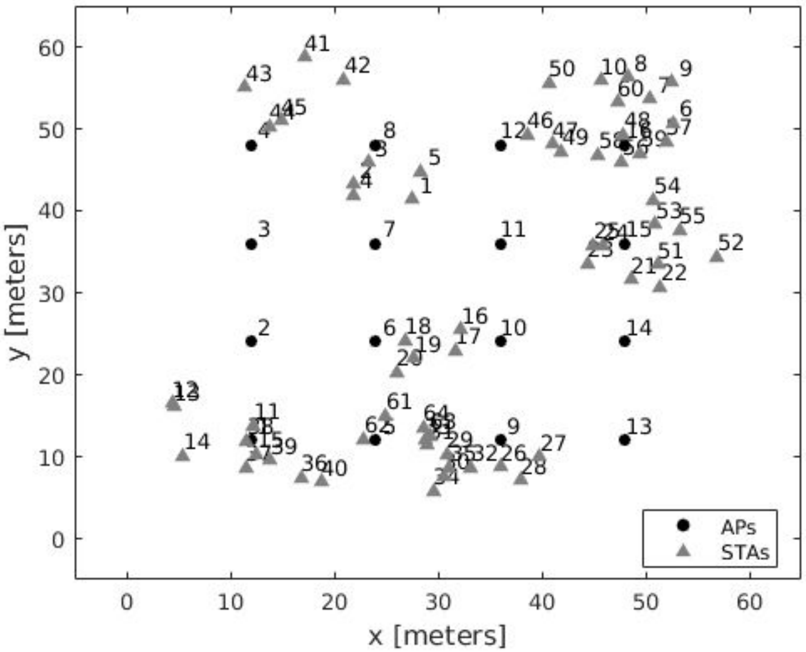

Cluster of 5 STAs

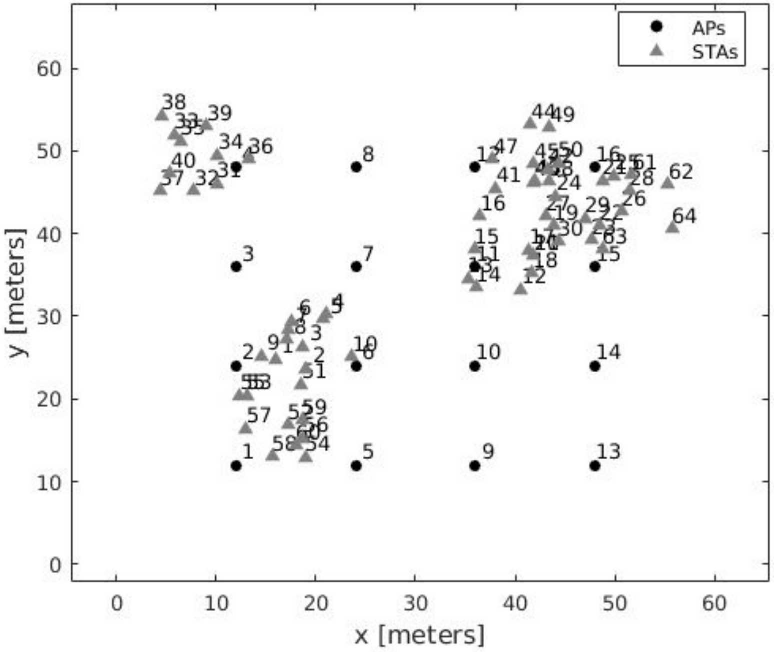

Cluster of 10 STAs

Standard association for WiFi

- STA checks all available channels
- Stores the signal strength for all APs
- Then selects the strongest signal for association
- Why is this the default method?
  - Signal strength defines data rate!
  - Lowest data rate* is 6.5 Mbps
  - Highest is 78 Mbps

* 802.11ac, single antenna, 20 MHz

Association for each case

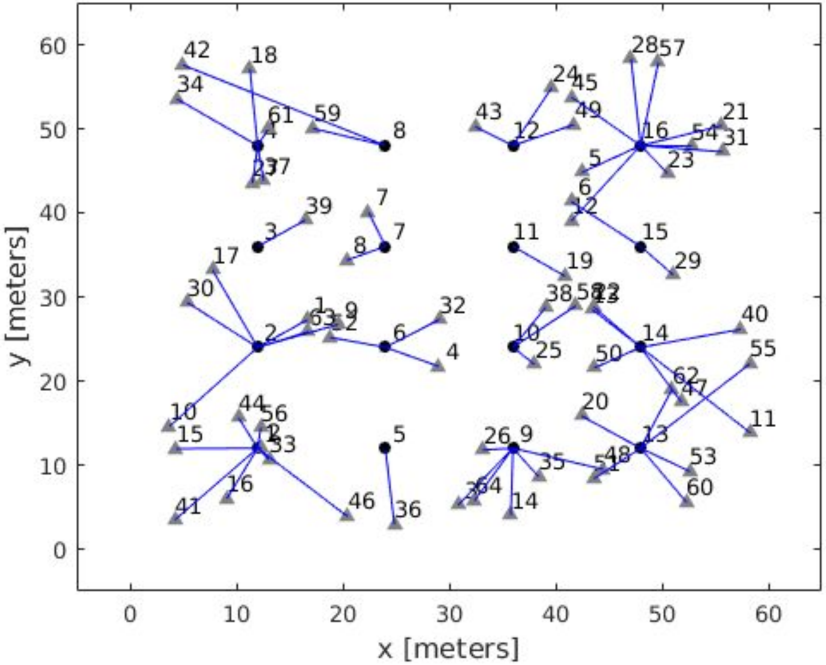

Uniform distribution of STAs

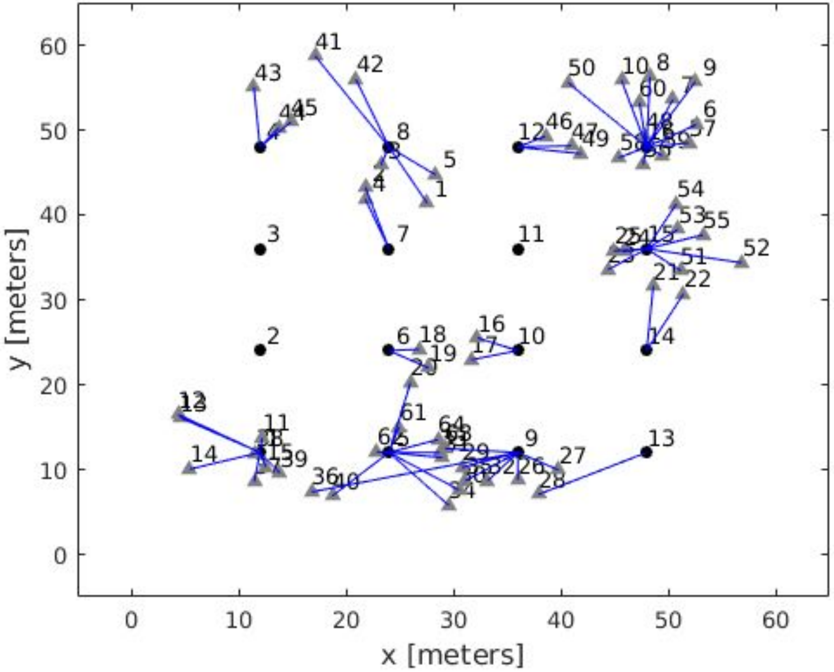

Clustered distribution of STAs

Airtime

Airtime is the time (per second) that a STA needs to transmit their load

For example, if we want to transmit 4 Mbps with packets of 12000 bits that require 1 ms to be transmitted, every second we need:

$$\frac{\mathrm{Load}}{\mathrm{L}} \cdot \mathrm{T} = \frac{4 \cdot 10^6}{12 \cdot 10^3} \cdot 0.001 = \boxed{333.33} \cdot 0.001 = 0.333$$

Packets/second

We need 0.333 seconds of airtime per second

Airtime

Let’s say that we have:

- STA 1 that needs 0.333 seconds
- STA 2 that needs 0.5 seconds
- STA 3 that needs 0.4 seconds

The AP then needs to account for 1.233 seconds of transmissions every second (which won’t work), meaning that we split the airtime proportionally:

- STA 1 gets 0.333/1.233 = 0.270 seconds
- STA 2 gets 0.5/1.233 = 0.405 seconds
- STA 3 gets 0.4/1.233 = 0.324 seconds

No one can transmit fully when the AP is saturated

- Execute file Main.m. You should get the following outputs:
  - A scatter plot showing the position of the AP and of the stations.
  - A bar plot showing the throughput of each STA.
  - A bar plot showing the satisfaction of each STA.
  - A console output with the STA and AP association for every node.
- Find where you can change the value of the random seed. Generate two new topologies. Discuss the observed results, explaining the observed differences.
- Find where you can change the number of stations. Try the following values: 5, 10, 20, 30, 40. Explain and justify the results regarding the throughput and satisfaction.
- If you could change the association manually, what would you do? Randomly? Would you call your solution 'intelligent'?
- For next session, read: `https://arxiv.org/pdf/1903.00281.pdf`

For next session

- Homework:
  - Read the following two papers:
    - Carrascosa Zamacois, Marc, and Boris Bellalta. **"Decentralized AP selection using multi-armed bandits: opportunistic ε-greedy with stickiness."** In *Proceedings of the 2019 IEEE Symposium on Computers and Communications (ISCC); 2019.*
    - Carrascosa, Marc, and Boris Bellalta. "**Multi-armed bandits for decentralized AP selection in enterprise WLANs**." *Computer Communications* 159 (2020): 108-123.