WS501 · WI-FI CSI SENSING FROM FIRST PRINCIPLES · WEEK 1

The Wireless Channel From a Signal to a Fingerprint

How a Wi-Fi signal bouncing around a room becomes a measurable fingerprint — and why that fingerprint changes when people move. No prior knowledge needed.

Jakub Dubec · jakub.dubec@stuba.sk · FIIT STU

This whole week in one sentence

A Wi-Fi signal bounces off the walls, floor, and people in a room. All those echoes add up into one number per frequency — the channel — and when a person moves, that number changes. That change is our sensor.

Everything else this week is detail under this one sentence.

Lecture 0 · First, the words

What this field is, the words you need, and why bother

What is Wi-Fi sensing?

One plain sentence.

THE FIELD IN A SENTENCE

Wi-Fi sensing uses the way radio waves bounce around a room to detect, count, or locate people — without a camera and without anyone carrying a device.

🏟️ **Shout in an empty gym** and you hear one clean echo. **Fill it with milling people** and the echoes blur into a constant, shifting mush. Radio does exactly the same thing — and this week is about how we **measure that mush.**

The words you need (1 / 2)

Channel — the single summary of everything that happens to a radio signal between the transmitter and the receiver: every echo, every bit of power lost on the way, rolled into one number (per frequency).

CSI — Channel State Information — the measured snapshot of that channel: how the signal arrived, across many tiny sub-channels at once. Think of it as the signal's **fingerprint** of the room.

Multipath — radio does not travel in one straight line. It bounces off walls, floor, and furniture and arrives by **many paths at once**. Each path is one echo.

LOS / NLOS — **line-of-sight** means a clear straight path exists between transmitter and receiver. **Non-line-of-sight** means that path is blocked, so only bounced echoes arrive.

The words you need (2 / 2)

dB — decibel — a unit for ratios of power, on a logarithmic scale. Rule of thumb: **+3 dB = double the power**, **−10 dB = one tenth** the power. It lets us add losses instead of multiplying them.

Phasor — a way to draw a wave as a little **rotating arrow**: its length is the wave's strength, its angle is the wave's phase (where it is in its cycle). Adding waves becomes adding arrows.

Fourier transform — a tool that asks "what mix of steady tones makes up this pattern?" It turns a pattern over **time** into a pattern over **frequency** — like splitting white light through a prism into colours.

OFDM — the Wi-Fi waveform that splits one channel into many narrow **sub-channels** (subcarriers) carried at once. Week 2 covers it; for now, just "Wi-Fi sends many tones side by side."

Why bother counting people?

SAFETY

Spot dangerous crowding **before** it becomes a crush.

COMFORT

Know which rooms are full and which are free.

ENERGY

Heat, cool, and ventilate **only** the rooms people are actually in.

Can we do it for free, with the Wi-Fi signals already in the air — no camera, no app?

> To even attempt that, we first have to understand what a Wi-Fi signal turns into on its way across a room. That is the whole job of Week 1.

Five things you will be able to do

1. **Explain how a Wi-Fi signal bouncing off walls becomes a mathematical fingerprint of the room — and build the formula for it from scratch.**
2. **Tell apart a steady link from a wandering one, and read off a single number (the K-factor) that says which it is.**
3. **Say how much power a plasterboard wall steals from a Wi-Fi signal at 2.4 GHz versus 5.0 GHz, using the lab's measured data.**
4. **Argue why the shape of a floorplan decides which measurements are even possible — using a real campaign that failed one of its own design goals.**
5. **Explain why a person walking through a room is detectable as a change in the channel — the bridge to next week.**

Part 1 · From bounces to a channel

Building the channel formula, one step at a time

A signal arrives many times over

- The transmitter sends out a wave at one instant
- A **direct** copy arrives first, by the short straight path
- **Reflected** copies bounce off walls, ceiling, floor — each takes longer and is weaker
- **Diffracted** copies bend around edges and corners
- The receiver hears **all the copies added together, at once**

> 🏔️ Like shouting in a canyon: you hear your own voice, then the cliff echo a moment later, then a fainter echo off the far wall — all overlapping.

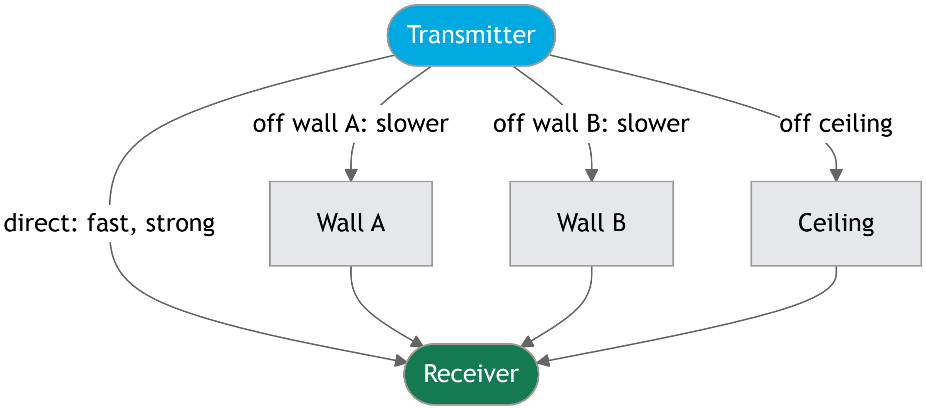

The first formula: adding up the echoes

Give each echo two numbers: how late it arrives, and how strong it is.

1. Call echo number *i* 's **delay** τ_i (Greek "tau") — how many seconds late it arrives.
2. Call its **strength** r_i — how much of the original wave survived that path.
3. The received signal is just every delayed, shrunk copy of the sent signal x, added up — plus a little background **noise** n.

$$y(t) = \sum_i r_i \, x(t - \tau_i) + n(t)$$

The big **Σ** means "add them all up." y is what we receive over time t; x is what was sent; each echo i is the sent wave delayed by τ_i and scaled by r_i. The last term n is thermal noise.

A picture before the next tool

🌈 **A Fourier transform** asks: what mix of steady tones makes up this pattern over time? Like an audio equaliser splitting a song into bass, mid, and treble — or a prism splitting white light into colours. We use it to turn our "echoes over time" into a "channel across frequencies."

Time-shift property — the one fact we borrow from the Fourier transform: a **delay in time** becomes a **steady rotation in frequency**. The later an echo, the faster its arrow spins as you sweep across frequency.

Delay in time → rotation in frequency. That is the whole trick on the next slide.

Building H(f) — step by step

Apply the Fourier transform to the echo sum. A delay becomes a phase rotation.

```
Step 5 — Add noise: the model we use all semester

  Y = H · X + N        (at one frequency / one subcarrier)

  H(f) = Σᵢ rᵢ · e^(-j2πf τᵢ)

  The channel H multiplies what we sent, then noise is added on top.
```

What that formula means: the channel is a number per frequency

$$H(f) = \sum_i r_i \, e^{-j2\pi f \tau_i}$$

At each frequency f, the channel H is the sum of all the echoes. Each echo is a rotating arrow (a phasor) of length r_i; its angle grows with frequency × that echo's delay τ_i.

> **Reminder — phasor** — a wave drawn as a rotating arrow: length = strength, angle = phase. Adding waves = adding arrows tip-to-tail.

And the system model, in plain words: what you receive = (the channel) × (what you sent) + noise.

Why the channel ripples across frequency

🎶 Picture several musicians playing the same note but starting at slightly different times. At some pitches they line up and swell; at others they fall out of step and cancel. Sweep across pitches and you hear a rippling loud-soft pattern — that is exactly what the channel does across frequency.

Figure 1 - A 3-path channel and the frequency-selective |H(f)|^2 it produces

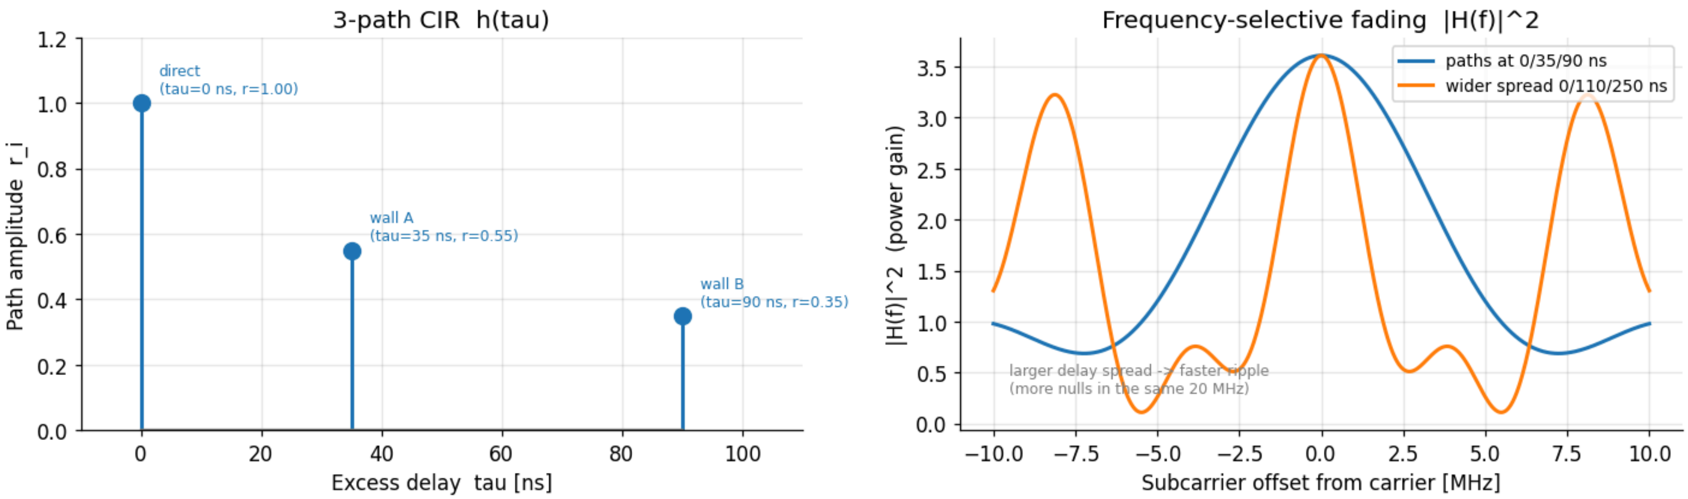

The channel is set entirely by the echoes — their strengths and delays

Move a person, open a door, and the echoes change. So the channel changes.

That is why the Wi-Fi channel quietly carries information about the room.

Part 2 · Steady link or wandering link?

Reading one number — the K-factor — off the channel

Two kinds of link — an everyday picture

A torch in a dark room: one strong, steady beam plus faint glints off the walls. The total stays roughly constant — you can rely on it. This is a **line-of-sight (LOS)** link.

Light through fog: no single dominant ray, just scattered glow arriving from everywhere. The brightness flickers and can briefly drop near zero. This is a **non-line-of-sight (NLOS)**, scatter-only link.

Fading — the channel strength going up and down. **Rayleigh fading** = the foggy case (no dominant ray). **Rician fading** = the torch case (one dominant ray plus scatter).

Rayleigh fading — the foggy, no-anchor case

Many comparable echoes, none dominant. The channel-arrow wanders randomly around the origin.

$$p(r) = \frac{r}{\sigma^2} e^{-r^2/2\sigma^2}, \qquad r \geq 0$$

This curve says how likely each channel strength r is. σ^2 (sigma-squared) is the average scatter power. The shape: moderate strengths are most common, but the channel can occasionally collapse to near zero — a deep fade.

> This describes a link cut off from any direct path — behind heavy obstruction, deep into a building. A link like this is unreliable on its own.

Rician fading — one dominant ray steadies the link

A strong direct ray, plus the usual scatter around it. The channel-arrow now wanders around a fixed point, not the origin.

> **K-factor** — the single number that captures this: how much power is in the one dominant ray versus all the scattered echoes.

$$K = \frac{\text{power in the dominant ray}}{\text{power in the scatter}}$$

Big K → a strong ray dominates → steady, reliable link (torch). K near zero → no ray wins → the channel behaves like Rayleigh (fog).

> In decibels: K(dB) = $10 \cdot \log_{10} K$. So K = 10 in power is +10 dB; K = 1 is 0 dB; K below 0 dB means scatter wins.

Measuring K from data — and the full curve (optional)

HOW WE READ K OFF REAL DATA

The **Greenstein moment estimator**: watch how jittery the channel *power* is over time. A steady link (torch) barely jitters → big K. A wandering link (fog) jitters a lot → K near zero. No curve-fitting needed.

> **Moment estimator** — a recipe that reads a quantity (here K) straight from simple summaries of the data: its average and its spread.

THE FULL RICIAN CURVE (FOR COMPLETENESS)

$$p(r) = \frac{r}{\sigma^2}\, e^{-(r^2+\nu^2)/2\sigma^2}\, I_0(\frac{r\nu}{\sigma^2})$$

ν is the dominant-ray strength; σ^2 is the scatter power. I_0 is a standard mathematical function — you never evaluate it by hand. As $\nu \to 0$ (dominant ray gone), this collapses back to the Rayleigh curve.

What the lab actually measured for K

Where the link runs	K (in dB)
Open room, clear line of sight	**10 – 15**
Through one plasterboard wall	weakly steady
Through the wall, scatter-dominated	**below 0** (fog-like)
Deep inside, no direct path	→ 0 (Rayleigh)

> A clear-LOS link is firmly "torch" (K = 10–15 dB). A through-wall link in this lab drifts toward "fog" (K below 0 dB) — because the scattered field wins once the direct ray is weakened.

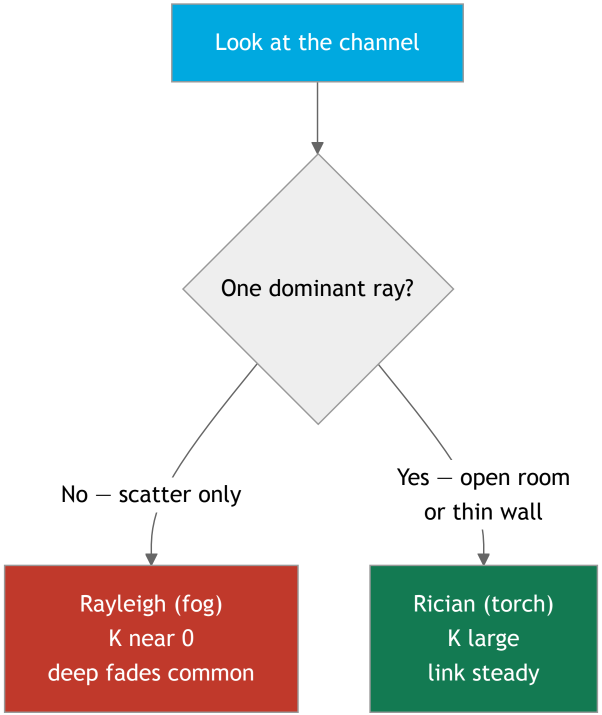

Recovering K from data — the figure

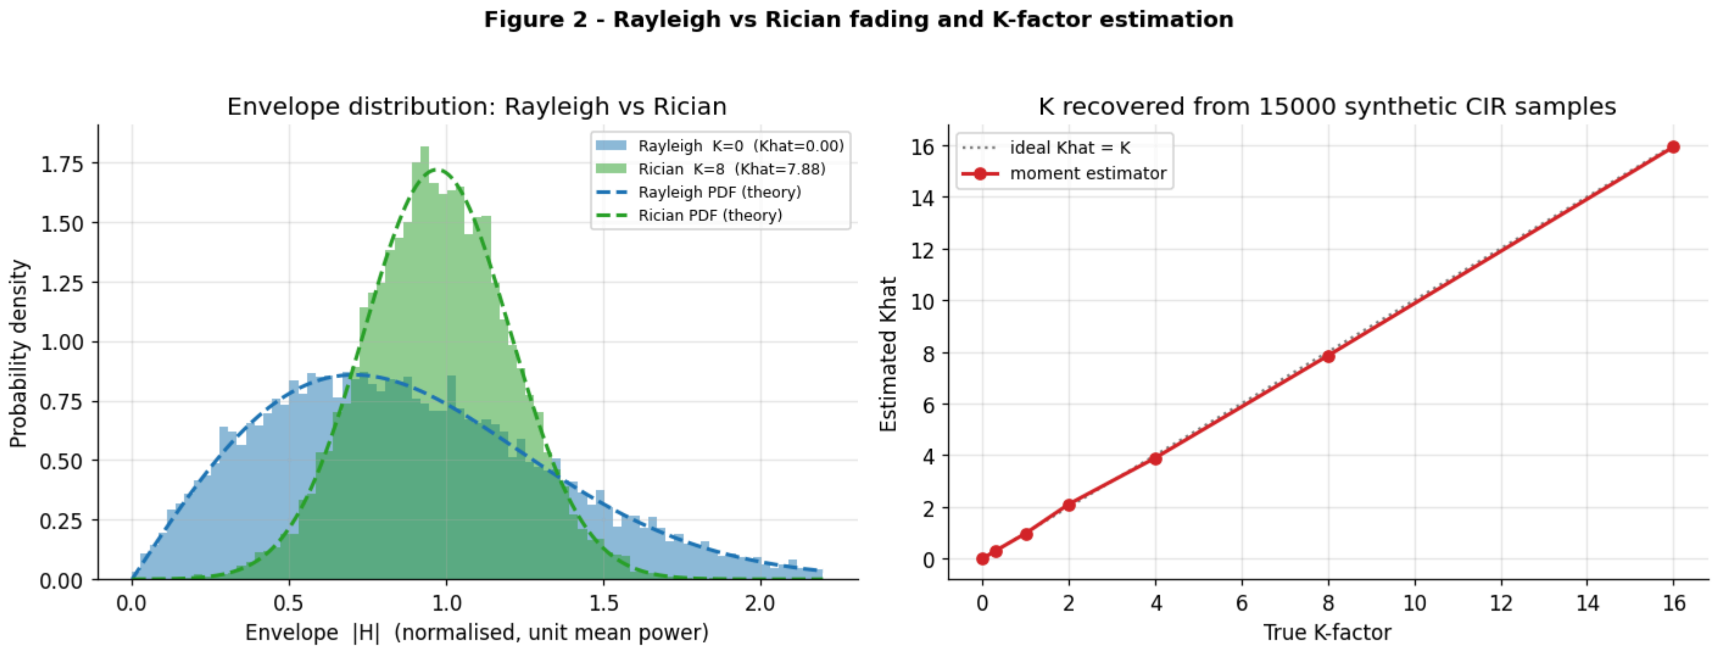

Left: simulated channel-strength histograms for a foggy (K = 0) and a torch-like (K = 8) link, with the theory curves overlaid. Right: the moment estimator recovers the true K across two decades — accurate to about **1%** over K = 0 to 16, from 15 000 simulated samples. Reproduce in workbook §2.

Part 3 · What a wall costs

How much power a plasterboard wall steals — measured

First, free space: power spreads out

☀️ A light bulb looks dimmer the farther you stand — its energy spreads over an ever-bigger sphere. Radio does the same: the farther the receiver, the smaller the slice of signal it catches.

dB — reminder — a ratio of power on a log scale: +3 dB doubles the power, −10 dB cuts it to a tenth. We measure every loss in dB so we can add them up.

ONE FACT FOR LATER

At a fixed distance, the higher band (5.0 GHz) starts out with roughly **6.5 dB more** free-space loss than the lower band (2.4 GHz), because a fixed-size antenna catches a smaller slice of a shorter wavelength.

The headline number: loss per wall

From the empty-room baseline (no people), each wall the signal crosses adds a fixed power penalty.

Band	Loss per plasterboard wall
2.4 GHz	**≈ 0.5 dB**
5.0 GHz	**11.0 dB**

At 2.4 GHz a plasterboard wall is **almost transparent** — it barely dims the signal. At 5.0 GHz the same wall is like frosted glass — it costs 11 dB, a roughly **22-fold** drop in power.

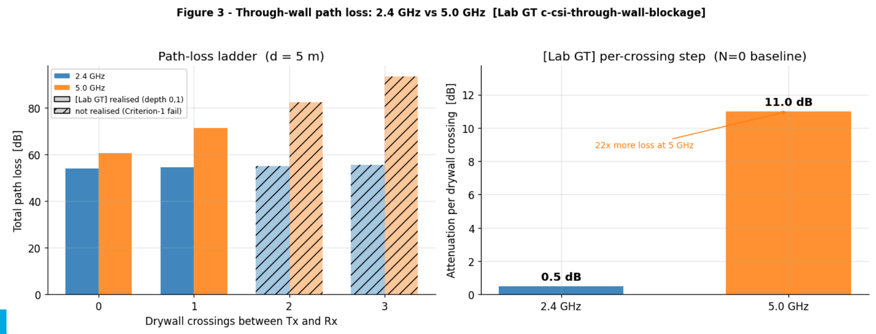

Total loss versus walls crossed, at 2.4 vs 5.0 GHz. The two-or-more-wall bars are hatched because the lab never managed to build such a link — the reason is Part 4. Reproduce in workbook §3.

Why the band choice can break an algorithm

THE TRAP

Some crowd-counting methods work by detecting the sharp **dip** when a person crosses the direct line. Those methods were validated in **concrete** buildings, where one wall removes 10–20 dB and the dip is obvious.

Part 4 · The floorplan is an experiment variable

A campaign that failed one of its own design goals

The lab’s through-wall campaign, in plain words

WHAT WAS MEASURED

- **52 measurement runs** on one test floor
- Crowd size swept from **0 to 6 people**
- Two bands: **2.4 and 5.0 GHz**
- Three repeats (seeds) per setting, for robustness

THE FIVE RADIO LINKS

- **One in-room link** — crosses no walls
- **Four through-wall links** — cross one wall each

Holding distance roughly fixed isolates the wall effect from plain free-space spreading.

> The empty-room runs give the per-wall loss (Part 3). The crowd runs give the per-person effect (a slide from now). But notice: only **two** kinds of link were ever physically built. Why? That is the next slide.

A design goal that failed

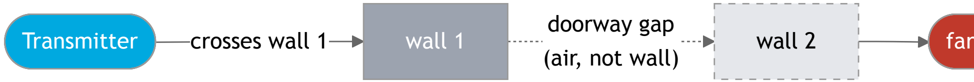

THE GOAL (SET IN ADVANCE)

At least one link should cross **two or more walls**, so we could measure the two-wall penalty and check the loss-per-wall pattern keeps going.

The number of walls a signal crosses depends on the exact shape of the floor — not on the room names.

Three lessons from a result that "failed"

1. **The floorplan decides what is measurable.** Which wall-counts are even possible depends on where you place transmitter and receiver in the actual room shapes. Shift a receiver 10 cm past a doorway and the wall-count — and 11 dB at 5 GHz — changes.

2. **Ground truth needs checking, not guessing.** "This link crosses two walls" guessed from a room map can be wrong. You have to trace the signal's path through the real wall outlines. The workbook does exactly this (§3) and recovers the true crossing counts.

3. **A negative result is still a measurement.** The shortfall was recorded, not buried, with a recommendation to move the far receiver off the doorway line before any re-test. Reporting only the conditions that worked would hide the limits of the method.

How a crowd shows up: the per-person effect

Each extra person in the room takes a little more power out of a through-wall link.

$$\frac{\text{change in signal}}{\text{per person}} = \begin{cases} -0.050 \pm 0.019\ \text{dB} & (2.4\ \text{GHz}) \\ -0.064 \pm 0.024\ \text{dB} & (5.0\ \text{GHz}) \end{cases}$$

Each added person costs the through-wall link about a twentieth of a dB at 2.4 GHz, a touch more at 5 GHz. The ± is the spread across repeats — small, noisy, but consistent.

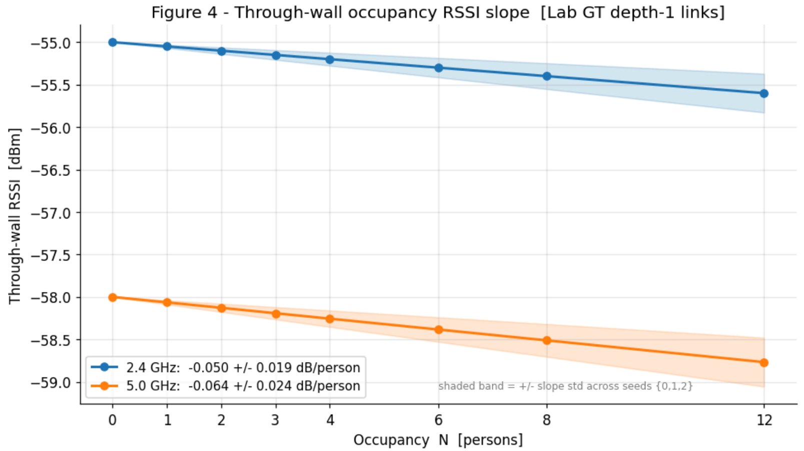

Part 5 · Why a moving person is visible

Coherence time and coherence bandwidth

Coherence bandwidth — an everyday picture

🎹 Press two **neighbouring** piano keys and they sound almost the same — they rise and fall together. Press two keys **far apart** and they behave independently. Subcarriers in Wi-Fi are the same: those close enough together fade together; those far apart fade on their own.

Coherence bandwidth (B_0) — how far apart in frequency two subcarriers must be before they fade independently. Closer than B_0 → they move together. Farther → they are on their own.

Delay spread (σ_τ) — how spread out in time the echoes are: the gap between the earliest and the stragglers.

Coherence bandwidth — the number

$$B_0 \approx \frac{1}{2\pi\,\sigma_\tau}$$

The more spread out the echoes (bigger σ_τ), the smaller the band B_0 over which two subcarriers fade together. Long echoes → fast ripple → narrow B_0.

A typical room (about 20 m across) has echoes spread over σ_τ ≈ 20–50 ns, which gives:

WORKED NUMBERS

B_0 ≈ **3 – 8 MHz**. A 20 MHz Wi-Fi channel is several B_0 wide — so it is **frequency-selective**: different parts of the band fade differently.

WHY THIS MATTERS FOR CSI

Subcarriers closer than B_0 are **not** independent — they fade in groups. With B_0 ≈ 6 MHz and Wi-Fi subcarriers spaced 312.5 kHz apart, about **19 neighbouring subcarriers** move together.

Recovering the echoes from the channel (and its limit)

Run the Fourier transform backwards (an IFFT) on the measured channel and you get the echoes back — but only coarsely.

$$\text{echoes over time} = \text{IFFT}\{\,\text{channel across frequency}\,\}$$

IFFT = inverse Fourier transform: frequency view → time view. This recovers the echo pattern (the channel impulse response).

THE RESOLUTION LIMIT

With a 20 MHz channel, the finest time-gap we can resolve is about **50 ns** — which is a path-length difference of **c · 50 ns ≈ 15 m**. Echoes closer than that blur into one. Coarse and aliased — Week 2 makes this precise for real hardware.

Coherence time — an everyday picture

🚑 An ambulance siren rises in pitch as it races toward you and drops as it passes — the **Doppler effect**. A person walking through the room shifts the radio echoes the same way; the faster they move, the bigger the shift, and the quicker the channel changes.

Coherence time (T_0) — how long the channel stays roughly the same before motion changes it. Fast motion → short coherence time.

Doppler shift — the change in a wave's frequency caused by something moving. For radio and a walker, it sets how fast the channel evolves.

Coherence time — the worked numbers

$$f_{\max} = \frac{v}{\lambda}, \qquad T_0 \approx \frac{1}{2 f_a}$$

f_max is the biggest Doppler shift, set by walking speed v over wavelength λ. The channel changes over the Doppler spread f_a ≈ 2·f_max, and T_0 is roughly its inverse.

WALKING AT 2.4 GHZ

```
v = 1.2 m/s      λ = 12.5 cm
f_max = v/λ  ≈ 9.6 Hz
f_a   ≈ 2·f_max ≈ 19 Hz
T_0   ≈ 1/(2·19) ≈ 26 ms
```

CAN WE KEEP UP?

A typical access point takes a fresh CSI snapshot every **10 ms** (100 times a second). The channel holds still for about **26 ms** — so we capture it several times before it changes. Comfortably fast enough.

So why is a slow walker so easy to detect?

THE PUZZLE

A person at 1.2 m/s moves only **4.8 cm** between two snapshots taken 10 ms apart. That is tiny. Why isn't it invisible?

THE ANSWER: PHASE IS A FINE RULER

Moving an echo's path by just **3 cm** (a quarter wavelength at 2.4 GHz) rotates its arrow by a quarter turn — a **big** change in the summed channel.

Detectable not because the person moves far — but because the channel is phase-sensitive.

> This same phase sensitivity is why CSI can later pick up **breathing** — chest motion of well under a centimetre.

The CIR figure and channel decorrelation

Figure 5 - CIR recovery by IFFT, coherence bandwidth and coherence time

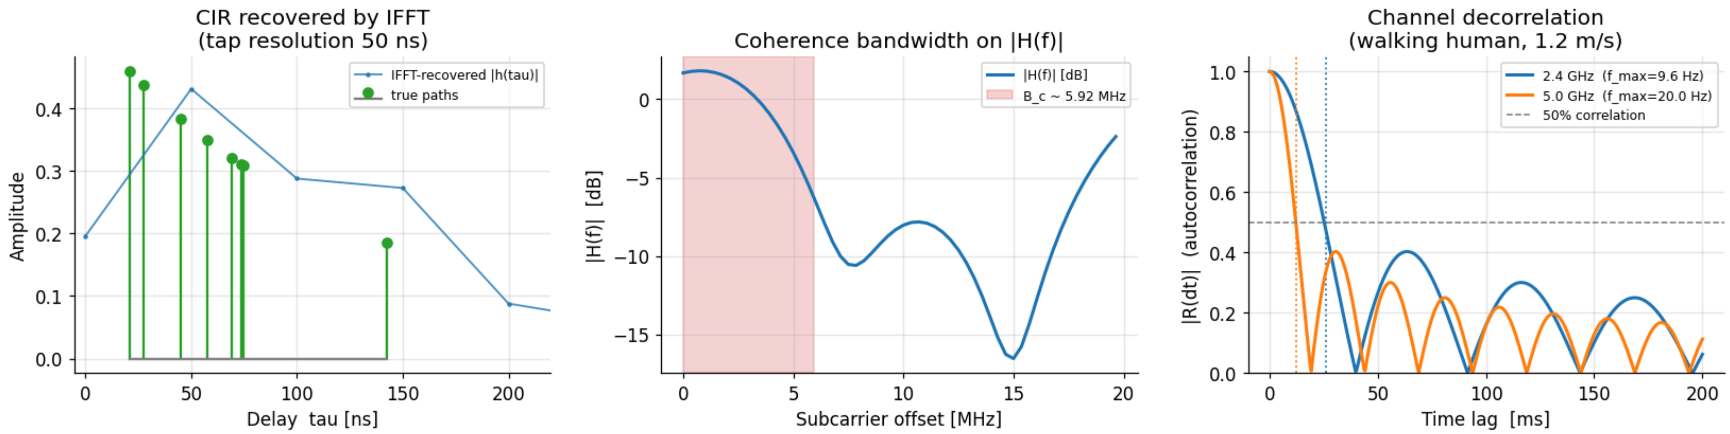

Left: the echo pattern recovered by inverse Fourier transform (blue) versus the true echoes (green) — the 50 ns grid blurs nearby ones. Centre: the coherence bandwidth shaded on the channel. Right: how fast the channel decorrelates as time passes for a walking person, with the coherence time marked on each band. Reproduce in workbook §5.

Bridge to Week 2: from the channel to the CSI matrix

WEEK 1 ✓

Built the channel from echoes. Steady vs wandering links, wall loss, coherence.

WEEK 2 →

How a Wi-Fi card actually samples the channel: OFDM, the preamble, the CSI matrix.

WEEK 3+

Turning CSI into answers: occupancy, activity, location.

> A receiver recovers the channel by sending a **known** block of symbols (a *preamble*) and dividing the received signal by what it knows it sent. The full story — OFDM structure, the errors that creep in — is Week 2.

Week 1 — recap, against the five goals

- **Goal 1 — the formula.** The channel = the sum of all echo-arrows: each has a strength and a delay. We built it step by step, not asserted it.
- **Goal 2 — steady vs wandering.** One number, the K-factor, says whether a dominant ray steadies the link (big K, "torch") or scatter wins (K near 0, "fog"). We read it off the data.
- **Goal 3 — what a wall costs.** Measured: **0.5 dB per wall at 2.4 GHz**, **11.0 dB at 5.0 GHz** — a 22-fold gap that decides which band to use.
- **Goal 4 — geometry is a variable.** A planned two-wall link didn't exist, because a doorway gap broke the second wall. The floor's shape, not the room labels, sets what you can measure.
- **Goal 5 — why motion shows.** The channel holds still for about **26 ms** while a person walks, and a 3 cm move swings it hugely — so a 100 Hz sampler catches it easily.

Try it yourself — the workbook

`workbook.ipynb` — five sections, five figures

§1 (Goal 1): build the channel formula; plot the channel strength for a 3-echo room and watch the dips appear.

§2 (Goal 2): simulate steady and wandering links at known K; recover K with the moment estimator; overlay the theory curves.

§3 (Goal 3): reproduce the loss-per-wall ladder; confirm the 0.5 and 11.0 dB steps.

§4 (Goal 4): plot the per-person effect (−0.050 / −0.064 dB) with its repeat-to-repeat spread.

§5 (Goal 5): recover the echoes by inverse transform; compute B_0 and T_0; show why motion is detectable.

Further reading

The sources behind this week

Further reading

- **Fallani et al. 2026** — derivation of the echo pattern and the channel in indoor multipath; how Wi-Fi cards acquire CSI from the preamble. *(fallani2026_04be)*
- **Ma et al. 2020** — broad survey of Wi-Fi sensing; free-space path loss; the Rician K-factor. *ACM Computing Surveys. (ma2020_4782)*
- **Wang et al. 2015** — multipath geometry for activity recognition; the path diagrams behind the echo-sum model. *MobiCom. (wang2015_48cf)*
- **Xie et al. 2015** — recovering the echo pattern from commodity hardware; coherence time and Doppler for walking people; resolution limits. *MobiCom. (xie2015_0389)*
- **Guarino et al. 2026** — the system model in element-wise form; reproducibility principles for sensing datasets. *(guarino2026_e72c)*
- **Depatla & Mostofi 2018** — counting people through walls with Wi-Fi; the line-crossing model, validated in concrete buildings — the contrast case for our plasterboard result. *PerCom. (depatla2018_97a9)*
- **Halperin et al. 2010** — the Linux 802.11n CSI Tool; how real hardware estimates the channel from the preamble. *SIGCOMM. (halperin2010_94f8)*
- **Zhang et al. 2024** — baseband channel formulation; carrier-frequency phase making the channel frequency-selective; indoor ray tracing. *ICNC. (zhang2024_df4a)*

Week 2: From OFDM to CSI — how a Wi-Fi receiver samples the channel on its subcarriers and turns it into the matrix the card reports.

Questions → course forum, tag #week1